HP StorageWorks

X1800sb and X3800sb Network Storage Blade

user guide

Part number: 572092-001
First edition: June 2009

Legal and notice information

© Copyright 1999, 2009 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

WARRANTY STATEMENT: To obtain a copy of the warranty for this product, see the warranty information website:

http://www.hp.com/go/storagewarranty

Intel, Itanium, Pentium, Intel Inside, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows XP, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Adobe and Acrobat are trademarks of Adobe Systems Incorporated.

Java is a US trademark of Sun Microsystems, Inc.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

UNIX is a registered trademark of The Open Group.

Printed in the US

Contents

About this guide 11
Intended audience 11
Related documentation 11
Document conventions and symbols 11
Rack stability 12
HP technical support 13
Customer self repair 13
Subscription service 13
HP websites 13
Documentation feedback 14

1 Administration tools 15
HP StorageWorks X1000 Automated Storage Manager 15
Microsoft Windows Storage Server 2008 administration tools 15
Remote Desktop for Administration 15
Share and Storage Management 16
Microsoft Services for Network File System 16
Single Instance Storage 17
Print Management 18

2 Storage management overview 19
Storage management elements 19
Storage management example 19
Physical storage elements 20
Arrays 21
Fault tolerance 22
Online spares 22
Logical storage elements 22
Logical drives (LUNs) 22
Partitions 23
Volumes 23
File system elements 24
File sharing elements 24
Volume Shadow Copy Service overview 24
Using storage elements 25
Clustered server elements 25
Network adapter teaming 25
Management tools 25
HP Systems Insight Manager 25
Management Agents 26

3 File server management 27
File services features in Windows Storage Server 2008 27
Storage Manager for SANs 27

Single Instance Storage 27
File Server Resource Manager 27
Windows SharePoint Services 27
File services management 28
Configuring data storage 28
Storage management utilities 28
Array management utilities 29
Array Configuration Utility 29
Disk Management utility 30
Guidelines for managing disks and volumes 30
Scheduling defragmentation 31
Disk quotas 31
Adding storage 32
Expanding storage 32
Extending storage using Windows Storage Utilities 33
Expanding storage for EVA arrays using Command View EVA 33
Expanding storage using the Array Configuration Utility 33
Volume shadow copies 34
Shadow copy planning 34
Identifying the volume 35
Allocating disk space 35
Identifying the storage area 36
Determining creation frequency 36
Shadow copies and drive defragmentation 37
Mounted drives 37
Managing shadow copies 37
The shadow copy cache file 38
Enabling and creating shadow copies 39
Viewing a list of shadow copies 40
Set schedules 40
Viewing shadow copy properties 40
Redirecting shadow copies to an alternate volume 41
Disabling shadow copies 41
Managing shadow copies from the storage system desktop 42
Shadow Copies for Shared Folders 42
SMB shadow copies 43
NFS shadow copies 44
Recovery of files or folders 45
Recovering a deleted file or folder 45
Recovering an overwritten or corrupted file 46
Recovering a folder 46
Backup and shadow copies 47
Shadow Copy Transport 47
Folder and share management 47
Folder management 48
Share management 54
Share considerations 54
Defining Access Control Lists 55
Integrating local file system security into Windows domain environments 55
Comparing administrative (hidden) and standard shares 55
Managing shares 56
File Server Resource Manager 56
Quota management 56
File screening management 57
Storage reports 57

Other Windows disk and data management tools 57
Additional information and references for file services 57
Backup 57
HP StorageWorks Library and Tape Tools 57
Antivirus 58

4 Cluster administration 59
Cluster overview 59
Cluster terms and components 60
Nodes 60
Resources 60
Cluster groups 61
Virtual servers 61
Failover and failback 61
Quorum disk 61
Cluster concepts 62
Sequence of events for cluster resources 62
Hierarchy of cluster resource components 63
Cluster planning 63
Storage planning 64
Network planning 64
Protocol planning 65
Preparing for cluster installation 66
Before beginning installation 66
Using multipath data paths for high availability 66
Checklists for cluster server installation 66
Network requirements 66
Shared disk requirements 67
Cluster installation 67
Setting up networks 68
Configuring the private network adapter 68
Configuring the public network adapter 68
Renaming the local area connection icons 68
Verifying connectivity and name resolution 68
Verifying domain membership 68
Setting up a cluster account 68
About the Quorum disk 68
Configuring shared disks 69
Verifying disk access and functionality 69
Configuring cluster service software 69
Using Cluster Administrator 69
Creating a cluster 69
Adding nodes to a cluster 70
Geographically dispersed clusters 70
Cluster groups and resources, including file shares 70
Cluster group overview 70
Node-based cluster groups 71
Load balancing 71
File share resource planning issues 71
Resource planning 71
Permissions and access rights on share resources 72
NFS cluster-specific issues 72
Non cluster aware file sharing protocols 73
Adding new storage to a cluster 73

Creating physical disk resources 73
Creating file share resources 73
Creating NFS share resources 74
Shadow copies in a cluster 74
Extend a LUN in a cluster 74
MSNFS administration on a server cluster 74
Best practices for running Server for NFS in a server cluster 75
Print services in a cluster 75
Creating a cluster printer spooler 75
Advanced cluster administration procedures 76
Failing over and failing back 76
Restarting one cluster node 77
Shutting down one cluster node 77
Powering down the cluster 77
Powering up the cluster 78

5 Troubleshooting, servicing, and maintenance 79
Troubleshooting the storage system 79
WEBES (Web Based Enterprise Services) 79
Maintenance and service 80
Maintenance updates 80
System updates 80
Firmware updates 80
Certificate of Authenticity 80

6 System recovery 81
The System Recovery DVD 81
To restore a factory image 81
The DON'T ERASE partition 81
Managing disks after a restoration 81

A Regulatory compliance notices 83
Regulatory compliance identification numbers 83
Federal Communications Commission notice 83
FCC rating label 83
Class A equipment 83
Class B equipment 84
Declaration of Conformity for products marked with the FCC logo, United States only 84
Modification 84
Cables 84
Canadian notice (Avis Canadien) 84
Class A equipment 84
Class B equipment 85
European Union notice 85
Japanese notices 85
Japanese VCCI-A notice 85
Japanese VCCI-B notice 85
Japanese power cord statement 85
Korean notices 86
Class A equipment 86
Class B equipment 86
Taiwanese notices 86
BSMI Class A notice 86

Taiwan battery recycle statement 86
Laser compliance notices 87
English laser notice 87
Dutch laser notice 87
French laser notice 88
German laser notice 88
Italian laser notice 88
Japanese laser notice 89
Spanish laser notice 89
Recycling notices 89
English notice 89
Bulgarian notice 90
Czech notice 90
Danish notice 90
Dutch notice 90
Estonian notice 91
Finnish notice 91
French notice 91
German notice 91
Greek notice 92
Hungarian notice 92
Italian notice 92
Latvian notice 92
Lithuanian notice 93
Polish notice 93
Portuguese notice 93
Romanian notice 93
Slovak notice 94
Spanish notice 94
Swedish notice 94
Turkish notice 94
Battery replacement notices 95
Dutch battery notice 95
French battery notice 96
German battery notice 96
Italian battery notice 97
Japanese battery notice 97
Spanish battery notice 98

Index 99

Figures

1 Storage management process example 20
2 Configuring arrays from physical drives 21
3 RAID 0 (data striping) (S1-S4) of data blocks (B1-B12) 21
4 Two arrays (A1, A2) and five logical drives (L1 through L5) spread over five physical drives 23
5 System administrator view of Shadow Copies for Shared Folders 38
6 Shadow copies stored on a source volume 38
7 Shadow copies stored on a separate volume 39
8 Accessing shadow copies from My Computer 42
9 Client GUI 44
10 Recovering a deleted file or folder 46
11 Properties dialog box, Security tab 49
12 Advanced Security settings dialog box, Permissions tab 50
13 User or group Permission Entry dialog box 51
14 Advanced Security Settings dialog box, Auditing tab 52
15 Select User or Group dialog box 52
16 Auditing Entry dialog box for folder name NTFS Test 53
17 Advanced Security Settings dialog box, Owner tab 54
18 Storage system cluster diagram 60
19 Cluster concepts diagram 62

Tables

1 Document conventions 11
2 Summary of RAID methods 22
3 Tasks and utilities needed for storage system configuration 28
4 Sharing protocol cluster support 65
5 Power sequencing for cluster installation 67

About this guide

This guide provides information about the following HP StorageWorks Network Storage Systems:

- HP StorageWorks X1800sb Network Storage Blade
- HP StorageWorks X3800sb Network Storage Gateway Blade

The HP StorageWorks X1800sb Network Storage Blade and HP StorageWorks X3800sb Network Storage Gateway Blade are based on HP ProLiant BL460c G6 Server Blade hardware. The HP StorageWorks X1800sb Network Storage Blade also can be used with an optional HP StorageWorks SB40c Storage Blade for internal shared storage and file serving or as a gateway to additional external storage.

For more information about these HP BladeSystem products, go to http://www.hp.com/go/blades.

Intended audience

This guide is intended for technical professionals with knowledge of:

- Microsoft® administrative procedures
- System and storage configurations

Related documentation

The following documents [and websites] provide related information:

- *HP StorageWorks X1800sb Network Storage Blade installation instructions*
- *HP StorageWorks X3800sb Network Storage Gateway Blade installation instructions*
- *HP StorageWorks X1000 Automated Storage Manager user guide*
- *HP Integrated Lights-Out 2 User Guide*

To obtain HP StorageWorks X1000 and X3000 Network Storage System documentation, go to http://www.hp.com/go/nas, select your product family, product model, click **Support for your product**, and then click **Manuals**.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text: Table 1	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses

Convention	Element
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

⚠ WARNING!

Indicates that failure to follow directions could result in bodily harm or death.

△ CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Rack stability

Rack stability protects personnel and equipment.

⚠ WARNING!

To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
- Ensure that the full weight of the rack rests on the leveling jacks.
- Install stabilizing feet on the rack.
- In multiple-rack installations, fasten racks together securely.
- Extend only one rack component at a time. Racks can become unstable if more than one component is extended.

HP technical support

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Customer self repair

HP customer self repair (CSR) programs allow you to repair your StorageWorks product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider, or see the CSR website:

http://www.hp.com/go/selfrepair

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com

- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals
- http://www.hp.com/support/downloads

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to storagedocsFeedback@hp.com. All submissions become the property of HP.

1 Administration tools

HP StorageWorks X1000 and X3000 Network Storage Systems include several administration tools to simplify storage system management tasks. HP StorageWorks X1000 Network Storage Systems include the HP X1000 Automated Storage Manager (ASM) in addition to HP storage utilities and Microsoft® Windows® Storage Server 2008 administration tools.

HP StorageWorks X1000 Automated Storage Manager

After installing and setting up your storage system, you can begin managing your storage using the HP Automated Storage Manager (HP ASM). HP ASM comes preinstalled on all HP X1000 Network Storage Systems.

NOTE:

HP ASM is not supported on HP X3000 Network Storage Systems.

ASM provides storage-allocation wizards that walk you through the process of allocating and configuring storage on your HP Network Storage System to host application data and shared folders. The storage-allocation wizards also allow you to schedule backups, snapshots, and replication of hosted application data and shared folders. Other wizards are provided to help you set up Exchange Server storage, SQL Server database storage, storage for user-defined applications, and storage for shared folders.

For more information about using the HP Automated Storage Manager, see the HP ASM online help or the *HP StorageWorks X1000 Automated Storage Manager user guide*. Go to http://www.hp.com/go/nas, select your product family, select your product model, click **Support for your product**, and then click **Manuals**.

Microsoft Windows Storage Server 2008 administration tools

Microsoft® Windows® Storage Server 2008 operating systems provide a user interface for initial server configuration, unified storage system management, simplified setup and management of storage and shared folders, and support for Microsoft iSCSI Software Target. It is specially tuned to provide optimal performance for network-attached storage. Windows Storage Server 2008 provides significant enhancements in share and storage management scenarios, as well as integration of storage system management components and functionality.

Remote Desktop for Administration

You can remotely administer storage systems by using Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode). You can use it to administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote Desktop for Administration is specifically designed for server management.

Remote Desktop for Administration does not require the purchase of special licenses for client computers that access the server. It is not necessary to install Terminal Server Licensing when using Remote Desktop for Administration.

You can use Remote Desktop for Administration to log on to the server remotely with any of the following features:

- Remote Desktop Connection
- Remote Web Administration
- Windows Server Remote Administration Applet

For more information, see the Windows Storage Server 2008 Help.

Share and Storage Management

With the Share and Storage Management snap-in provided in this release, you can more easily set up and manage shared folders and storage. Share and Storage Management provides the following:

- MMC-based management of shared folders and storage.
- Provision Storage Wizard for creating and configuring storage for file sharing and block sharing, including creating LUNs on storage subsystems, as well as creating and formatting volumes on LUNs or server disks.

NOTE:

You must have a VDS Hardware Provider that is appropriate for your storage system installed in order to provision storage on an iSCSI target. If you have Microsoft iSCSI Software Target running on a Windows Storage Server 2008 storage system, install the Microsoft iSCSI Software Target VDS Hardware Provider on the client computer.

- Provision a Shared Folder Wizard for creating and configuring shared folders that can be accessed by using either the server message block (SMB) or NFS protocol.
- Single Instance Storage (SIS) can be enabled or disabled for each volume that is displayed in Share and Storage Management. SIS recovers disk space by reducing the amount of redundant data stored on a volume. It identifies identical files, storing only a single copy of the file in the SIS Common Store, and replacing the files with pointers to the file in the SIS Common Store.

The Share and Storage Management snap-in makes it possible to complete most of the administrative tasks that are required to create and manage shared folders and volumes without having to use the Shared Folder Management, Storage Manager for SANs, or Disk Management snap-ins. These tasks include configuring quotas to restrict the quantity of data, configuring file screening to prevent certain file types or only allowing certain file types defined by the administrator, and enabling indexing.

For more information, see the Windows Storage Server 2008 Help.

Microsoft Services for Network File System

Microsoft Services for Network File System (NFS) is a component of Windows Storage Server 2008 that provides a file-sharing solution for enterprises that have a mixed Windows and UNIX environment. By using Microsoft Services for NFS, you can configure storage services to make it possible for users to store and access files on the storage system, and to transfer files between the storage system and UNIX computers by using the NFS protocol.

In Windows Storage Server 2008, you can complete the initial configuration of Microsoft Services for NFS by using the Microsoft Services for NFS Configuration Guide, which includes an Identity

Mapping Setup Wizard. After completing the initial Microsoft Services for NFS configuration, you can use the Microsoft Services for NFS console for ongoing maintenance and administration.

If you are using Windows Storage Server 2008 in an environment that does not include an Active Directory directory service domain, you can install Active Directory Application Mode (ADAM) on your storage system to provide user name mapping for NFS. Microsoft Services for NFS can use any RFC 2307 compliant Lightweight Directory Access Protocol (LDAP) service to provide username mapping services.

For more information, see the Windows Storage Server 2008 Help.

Single Instance Storage

The Single Instance Storage (SIS) feature reduces the amount of space that is used to store data on a volume. SIS does this by replacing duplicate files with logical links that point to a single copy of the file in the SIS Common Store, which is a hidden folder that is located in the root directory of the volume.

SIS consists of two primary components that together maintain a database of file signatures. These components include:

- Groveler service - The Groveler service scans the hard-disk volumes on a server for duplicate copies of files. If the service locates duplicate copies of files, the information about the duplicates is sent to the Single Instance Storage Filter. The Groveler service runs as a user-level service.
- Single Instance Storage Filter - The Single Instance Storage Filter is a file system filter service that manages duplicate copies of files on hard-disk volumes. When notified by the Groveler service of duplicate copies of files, this component copies one instance of a duplicate file into a central folder. The duplicate is then replaced by a link (a reparse point) to the central copy. The link file contains information about the original file, such as its current location, size, and attributes. The Single Instance Storage Filter runs in kernel mode.

The Single Instance Storage Filter service cannot be stopped. If this service is disabled, the linked files are not accessible. If the central folder is deleted, the linked files can become permanently inaccessible. If you stop the Groveler service, the files cannot be automatically linked, but the existing linked files can still be accessible.

You can enable SIS on a maximum of 20 volumes per computer. SIS cannot act upon any files that are referenced through junction points, and it cannot be used with any file system except the NTFS file system. SIS will not process files that are 32 kilobytes or less in size.

If you need to access data that is stored on a SIS volume, which might be required for backup and recovery operations, you must either run or have installed Single Instance Storage Filter on your computer.

Backup and recovery by using SIS has the following requirements:

- The backup software used must support SIS-enabled volumes.
- The SIS volume, SIS Common Store folder, and reparse points (links) to the files must be restored to a Windows 2000 NTFS version 5.0 (or later) file system or partition that supports reparse points or junction points.
- The Single Instance Storage Filter must be installed or enabled to access the data in the SIS volume.
- The backup program must be capable and configured to backup and restore the reparse points or junction points (links) to the files, and the SIS volume and the SIS Common Store folder must be selected.

For more information, see the Windows Storage Server 2008 Help.

Print Management

Print Management is an MMC snap-in that you can use to view and manage printers and print servers in your organization. You can use Print Management from any computer running Windows Storage Server 2008, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, Windows Storage Server 2003, Windows Storage Server 2003 R2, or Windows Storage Server 2008.

Print Management provides details such as the queue status, printer name, driver name, and server name. You can also set custom views by using the Print Management filtering capability. For example, you can create a view that displays only printers in a particular error state. You can also configure Print Management to send e-mail notifications or run scripts when a printer or print server needs attention. The filtering capability also allows you to bulk edit print jobs, such as canceling all print jobs at once. You can also delete multiple printers at the same time.

Administrators can install printers remotely by using the automatic detection feature, which finds and installs printers on the local subnet to the local print server. Administrators can log on remotely to a server at a branch location, and then install printers remotely.

For more information, see the Windows Storage Server 2008 Help.

2 Storage management overview

This chapter provides an overview of some of the components that make up the storage structure of the storage system.

Storage management elements

Storage is divided into four major divisions:

- Physical storage elements
- Logical storage elements
- File system elements
- File sharing elements

Each of these elements is composed of the previous level's elements.

Storage management example

Figure 1 depicts many of the storage elements that one would find on a storage device. The following sections provide an overview of the storage elements.

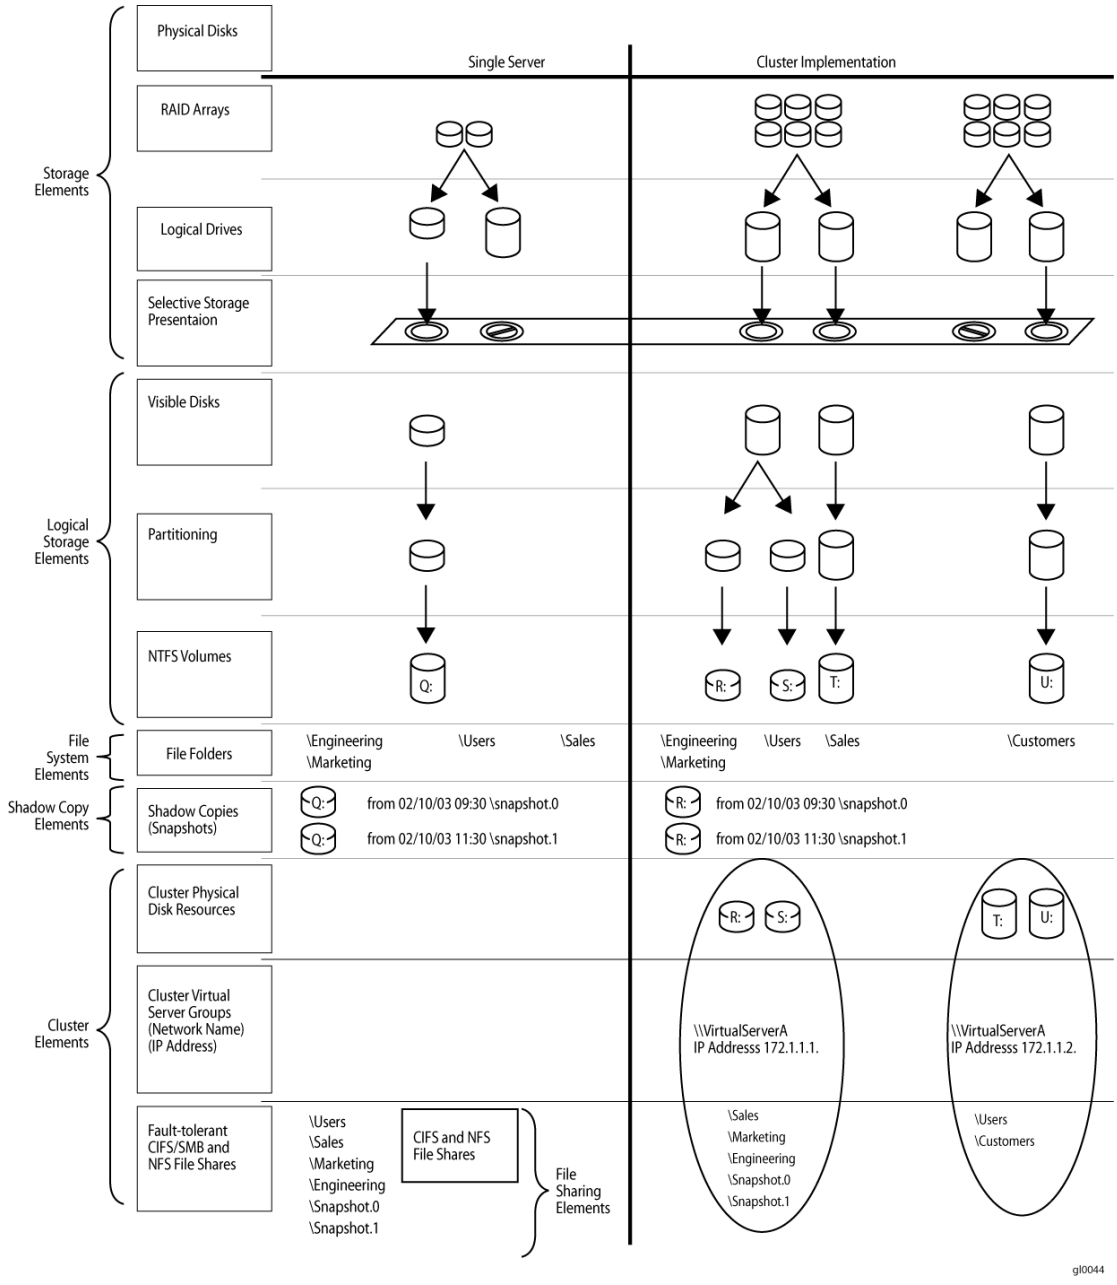

Figure 1 Storage management process example

Physical storage elements

The lowest level of storage management occurs at the physical drive level. Minimally, choosing the best disk carving strategy includes the following policies:

- Analyze current corporate and departmental structure.
- Analyze the current file server structure and environment.
- Plan properly to ensure the best configuration and use of storage.
  - Determine the desired priority of fault tolerance, performance, and storage capacity.
  - Use the determined priority of system characteristics to determine the optimal striping policy and RAID level.

- Include the appropriate number of physical drives in the arrays to create logical storage elements of desired sizes.

Arrays

See Figure 2. With an array controller installed in the system, the capacity of several physical drives (P1–P3) can be logically combined into one or more logical units (L1) called arrays. When this is done, the read/write heads of all the constituent physical drives are active simultaneously, dramatically reducing the overall time required for data transfer.

NOTE:

Depending on the storage system model, array configuration may not be possible or necessary.

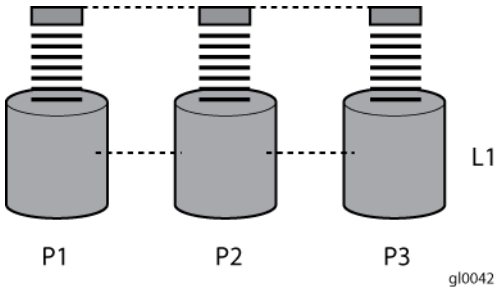

Figure 2 Configuring arrays from physical drives

Because the read/write heads are simultaneously active, the same amount of data is written to each drive during any given time interval. Each unit of data is termed a block. The blocks form a set of data stripes over all the hard drives in an array, as shown in Figure 3.

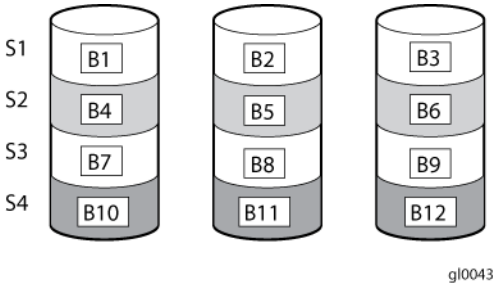

Figure 3 RAID 0 (data striping) (S1-S4) of data blocks (B1-B12)

For data in the array to be readable, the data block sequence within each stripe must be the same. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each hard drive in a given array contains the same number of data blocks.

NOTE:

If one hard drive has a larger capacity than other hard drives in the same array, the extra capacity is wasted because it cannot be used by the array.

Fault tolerance

Drive failure, although rare, is potentially catastrophic. For example, using simple striping as shown in Figure 3, failure of any hard drive leads to failure of all logical drives in the same array, and hence to data loss.

To protect against data loss from hard drive failure, storage systems should be configured with fault tolerance. HP recommends adhering to RAID 5 configurations.

The table below summarizes the important features of the different kinds of RAID supported by the Smart Array controllers. The decision chart in the following table can help determine which option is best for different situations.

Table 2 Summary of RAID methods

	RAID 0 Striping (no fault tolerance)	RAID 1+0 Mirroring	RAID 5 Distributed Data Guarding	RAID 6 (ADG)
Maximum number of hard drives	N/A	N/A	14	Storage system dependent
Tolerant of single hard drive failure?	No	Yes	Yes	Yes
Tolerant of multiple simultaneous hard drive failures?	No	If the failed drives are not mirrored to each other	No	Yes (two drives can fail)

Online spares

Further protection against data loss can be achieved by assigning an online spare (or hot spare) to any configuration except RAID 0. This hard drive contains no data and is contained within the same storage subsystem as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID level fault tolerance protection. However, unless RAID Advanced Data Guarding (ADG) is being used, which can support two drive failures in an array, in the unlikely event that a third drive in the array should fail while data is being rewritten to the spare, the logical drive still fails.

Logical storage elements

Logical storage elements consist of those components that translate the physical storage elements to file system elements. The storage system uses the Window Disk Management utility to manage the various types of disks presented to the file system. There are two types of LUN presentation: basic disk and dynamic disk. Each of these types of disk has special features that enable different types of management.

Logical drives (LUNs)

While an array is a physical grouping of hard drives, a logical drive consists of components that translate physical storage elements into file system elements.

It is important to note that a LUN may span all physical drives within a storage controller subsystem, but cannot span multiple storage controller subsystems.

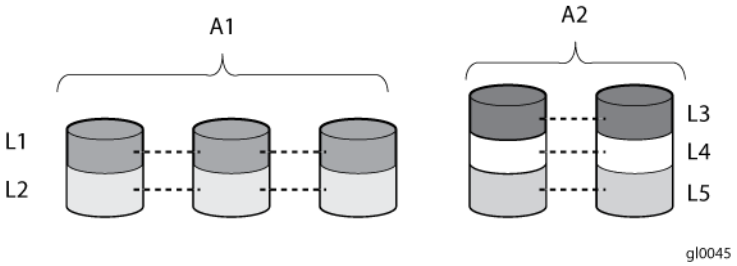

Figure 4 Two arrays (A1, A2) and five logical drives (L1 through L5) spread over five physical drives

NOTE:

This type of configuration may not apply to all storage systems and serves only as an example.

Through the use of basic disks, you can create primary partitions or extended partitions. Partitions can only encompass one LUN. Through the use of dynamic disks, you can create volumes that span multiple LUNs. You can use the Windows Disk Management utility to convert disks to dynamic and back to basic and to manage the volumes residing on dynamic disks. Other options include the ability to delete, extend, mirror, and repair these elements.

Partitions

Partitions exist as either primary partitions or extended partitions and can be composed of only one basic disk no larger than 2 TB. Basic disks can also only contain up to four primary partitions, or three primary partitions and one extended partition. In addition, the partitions on them cannot be extended beyond the limits of a single LUN. Extended partitions allow the user to create multiple logical drives. These partitions or logical disks can be assigned drive letters or be used as mount points on existing disks. If mount points are used, it should be noted that Services for UNIX (SFU) does not support mount points at this time. The use of mount points in conjunction with NFS shares is not supported.

Volumes

When planning dynamic disks and volumes, there is a limit to the amount of growth a single volume can undergo. Volumes are limited in size and can have no more than 32 separate LUNs, with each LUN not exceeding 2 terabytes (TB), and volumes totaling no more than 64 TB of disk space.

The RAID level of the LUNs included in a volume must be considered. All of the units that make up a volume should have the same high-availability characteristics. In other words, the units should all be of the same RAID level. For example, it would not be a good practice to include both a RAID 1+0 and a RAID 5 array in the same volume set. By keeping all the units the same, the entire volume retains the same performance and high-availability characteristics, making managing and maintaining the volume much easier. If a dynamic disk goes offline, the entire volume dependent on the one or more dynamic disks is unavailable. There could be a potential for data loss depending on the nature of the failed LUN.

Volumes are created out of the dynamic disks, and can be expanded on the fly to extend over multiple dynamic disks if they are spanned volumes. However, after a type of volume is selected, it cannot be altered. For example, a spanning volume cannot be altered to a mirrored volume without deleting and recreating the volume, unless it is a simple volume. Simple volumes can be mirrored or converted to spanned volumes. Fault-tolerant disks cannot be extended. Therefore, selection of the volume type

is important. The same performance characteristics on numbers of reads and writes apply when using fault-tolerant configurations, as is the case with controller-based RAID. These volumes can also be assigned drive letters or be mounted as mount points off existing drive letters.

The administrator should carefully consider how the volumes will be carved up and what groups or applications will be using them. For example, putting several storage-intensive applications or groups into the same dynamic disk set would not be efficient. These applications or groups would be better served by being divided up into separate dynamic disks, which could then grow as their space requirements increased, within the allowable growth limits.

NOTE:

Dynamic disks cannot be used for clustering configurations because Microsoft Cluster only supports basic disks.

File system elements

File system elements are composed of the folders and subfolders that are created under each logical storage element (partitions, logical disks, and volumes). Folders are used to further subdivide the available file system, providing another level of granularity for management of the information space. Each of these folders can contain separate permissions and share names that can be used for network access. Folders can be created for individual users, groups, projects, and so on.

File sharing elements

The storage system supports several file sharing protocols, including Distributed File System (DFS), Network File System (NFS), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), and Microsoft Server Message Block (SMB). On each folder or logical storage element, different file sharing protocols can be enabled using specific network names for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Volume Shadow Copy Service overview

The Volume Shadow Copy Service (VSS) provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. VSS supports 64 shadow copies per volume.

Shadow Copies of Shared Folders resides within this infrastructure, and helps alleviate data loss by creating shadow copies of files or folders that are stored on network file shares at pre-determined time intervals. In essence, a shadow copy is a previous version of the file or folder at a specific point in time.

By using shadow copies, a storage system can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer.

Shadow copies should not replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. For example, shadow copies cannot protect against data loss due to media failures; however, recovering data from shadow copies can reduce the number of times needed to restore data from tape.

Using storage elements

The last step in creating the element is determining its drive letter or mount point and formatting the element. Each element created can exist as a drive letter, assuming one is available, and/or as mount points on an existing folder or drive letter. Either method is supported. However, mount points cannot be used for shares that will be shared using Microsoft Services for Unix. They can be set up with both but the use of the mount point in conjunction with NFS shares causes instability with the NFS shares.

Formats consist of NTFS, FAT32, and FAT. All three types can be used on the storage system. However, VSS can only use volumes that are NTFS formatted. Also, quota management is possible only on NTFS.

Clustered server elements

HP StorageWorks X3000 Network Storage Systems support clustering. These storage systems support several file sharing protocols including DFS, NFS, FTP, HTTP, and Microsoft SMB. Only NFS, FTP, and Microsoft SMB are cluster-aware protocols. HTTP can be installed on each node but the protocols cannot be set up through cluster administrator, and they will not fail over during a node failure.

Network names and IP address resources for the clustered file share resource can also be established for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Network adapter teaming

Network adapter teaming is software-based technology used to increase a server's network availability and performance. Teaming enables the logical grouping of physical adapters in the same server (regardless of whether they are embedded devices or Peripheral Component Interconnect (PCI) adapters) into a virtual adapter. This virtual adapter is seen by the network and server-resident network-aware applications as a single network connection.

Management tools

HP Systems Insight Manager

HP SIM is a web-based application that allows system administrators to accomplish normal administrative tasks from any remote location, using a web browser. HP SIM provides device management capabilities that consolidate and integrate management data from HP and third-party devices.

IMPORTANT:

You must install and use HP SIM to benefit from the Pre-Failure Warranty for processors, SAS and SCSI hard drives, and memory modules.

For additional information, refer to the Management CD in the HP ProLiant Essentials Foundation Pack or the HP SIM website (http://www.hp.com/go/hpsim).

Management Agents

Management Agents provide the information to enable fault, performance, and configuration management. The agents allow easy manageability of the server through HP SIM software, and thirdparty SNMP management platforms. Management Agents are installed with every SmartStart assisted installation or can be installed through the HP PSP. The Systems Management homepage provides status and direct access to in-depth subsystem information by accessing data reported through the Management Agents. For additional information, refer to the Management CD in the HP ProLiant Essentials Foundation Pack or the HP website (http://www.hp.com/servers/manage).

3 File server management

This chapter begins by identifying file services in Windows Storage Server 2008. The remainder of the chapter describes the many tasks and utilities that play a role in file server management.

File services features in Windows Storage Server 2008

Storage Manager for SANs

The Storage Manager for SANs (also called Simple SAN) snap-in enables you to create and manage the LUNs that are used to allocate space on storage arrays. Storage Manager for SANs can be used on SANs that support Virtual Disk Server (VDS). It can be used in both Fibre Channel and iSCSI environments.

For more information on Storage Manager for SANs, see the online help.

Single Instance Storage

Single Instance Storage (SIS) provides a copy-on-write link between multiple files. Disk space is recovered by reducing the amount of redundant data stored on a server. If a user has two files sharing disk storage by using SIS, and someone modifies one of the files, users of the other files do not see the changes. The underlying shared disk storage that backs SIS links is maintained by the system and is only deleted if all the SIS links pointing to it are deleted. SIS automatically determines that two or more files have the same content and links them together.

File Server Resource Manager

File Server Resource Manager is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports.

By using File Server Resource Manager, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and to generate notifications when the quota limits are approached and exceeded.
- Create file screens to screen the files that users can save on volumes and in folders and to send notifications when users attempt to save blocked files.
- Schedule periodic storage reports that allow users to identify trends in disk usage and to monitor attempts to save unauthorized files, or generate the reports on demand.

Windows SharePoint Services

Windows SharePoint Services is an integrated set of collaboration and communication services designed to connect people, information, processes, and systems, within and beyond the organization firewall.

File services management

Information about the storage system in a SAN environment is provided in the HP StorageWorks SAN Manuals page located on the HP web site at www.hp.com/go/SDGManuals.

Configuring data storage

HP StorageWorks X1000 and X3000 Network Storage Systems are configured only for the operating system. The administrator must configure data storage for the storage system.

Configuring additional data storage involves creating arrays, logical disks, and volumes. Table 3 shows the general task areas to be performed as well as the utilities needed to configure storage for an HP Smart Array-based storage system.

Table 3 Tasks and utilities needed for storage system configuration

Task	Storage management utility
Create disk arrays	HP Array Configuration Utility or HP Automated Storage Manager
Create logical disks from the array space	HP Array Configuration Utility or HP Automated Storage Manager
Verify newly created logical disks	Windows Disk Management
Create a volume on the new logical disk	Windows Disk Management

- Create disk arrays—On storage systems with configurable storage, physical disks can be arranged as RAID arrays for fault tolerance and enhanced performance, and then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the storage system.

△ CAUTION:

The first two logical drives are configured for the storage system operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.

The fault tolerance level depends on the amount of disks selected when the array was created. A minimum of two disks is required for RAID 0+1 configuration, three disks for a RAID 5 configuration, and four disks for a RAID 6 (ADG) configuration.

- Create logical disks from the array space—Select the desired fault tolerance, stripe size, and size of the logical disk.
- Verify newly created logical disks—Verify that disks matching the newly created sizes are displayed.
- Create a volume on the new logical disk—Select a drive letter and enter a volume label, volume size, allocation unit size, and mount point (if desired).

Storage management utilities

The storage management utilities preinstalled on the storage system include the HP Array Configuration Utility (ACU).

Array management utilities

Storage devices for RAID arrays and LUNs are created and managed using the array management utilities mentioned previously. For HP Smart Arrays use the ACU.

NOTE:

The ACU is used to configure and manage array-based storage. Software RAID-based storage systems use Microsoft Disk Manager to manage storage. You need administrator or root privileges to run the ACU.

Array Configuration Utility

The HP ACU supports the Smart Array controllers and hard drives installed on the storage system.

To open the ACU from the storage system desktop:

NOTE:

If this is the first time that the ACU is being run, you will be prompted to select the Execution Mode for ACU. Selecting Local Application Mode allows you to run the ACU from a Remote Desktop, remote console, or storage system web access mode. Remote service mode allows you to access the ACU from a remote browser.

1. Select **Start > Programs > HP Management Tools > Array Configuration Utility**.
2. If the Execution Mode for ACU is set to Remote Mode, log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.

To open the ACU in browser mode:

NOTE:

Confirm that the ACU Execution Mode is set to remote service.

1. Open a browser and enter the server name or IP address of the destination server. For example, `http://servername:2301` or `http://192.0.0.1:2301`.
2. Log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.
3. Click **Array Configuration Utility** on the left side of the window. The ACU opens and identifies the controllers that are connected to the system.

Some ACU guidelines to consider:

- Do not modify the first two logical drives of the storage system; they are configured for the storage system operating system.
- Spanning more than 14 disks with a RAID 5 volume is not recommended.
- Designate spares for RAID sets to provide greater protection against failures.
- RAID sets cannot span controllers.
- A single array can contain multiple logical drives of varying RAID settings.

- Extending and expanding arrays and logical drives is supported.

The *HP Array Configuration Utility User Guide* is available for download at http://www.hp.com/support/manuals.

Disk Management utility

The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. Disk Management is used to initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Most disk-related tasks can be performed in Disk Management without restarting the system or interrupting users. Most configuration changes take effect immediately. A complete online help facility is provided with the Disk Management utility for assistance in using the product.

NOTE:

- When the Disk Management utility is accessed through a Remote Desktop connection, this connection can only be used to manage disks and volumes on the server. Using the Remote Desktop connection for other operations during an open session closes the session.
- When closing Disk Management through a Remote Desktop connection, it may take a few moments for the remote session to log off.

Guidelines for managing disks and volumes

- The first two logical drives are configured for the storage system operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.
- HP does not recommend spanning array controllers with dynamic volumes. The use of software RAID-based dynamic volumes is not recommended. Use the array controller instead; it is more efficient.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. (For example, volume e: might be named "Disk E:.") Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case the system needs to be restored.
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic, but cannot be converted back to basic without deleting all data on the disk.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of shadow copies, performance, and defragmentation.
- NTFS formatted drives are recommended because they provide the greatest level of support for shadow copies, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32.
- Read the online Disk Management help found in the utility.

Scheduling defragmentation

Defragmentation is the process of analyzing local volumes and consolidating fragmented files and folders so that each occupies a single, contiguous space on the volume. This improves file system performance. Because defragmentation consolidates files and folders, it also consolidates the free space on a volume. This reduces the likelihood that new files will be fragmented.

Defragmentation for a volume can be scheduled to occur automatically at convenient times. Defragmentation can also be done once, or on a recurring basis.

NOTE:

Scheduling defragmentation to run no later than a specific time prevents the defragmentation process from running later than that time. If the defragmentation process is running when the time is reached, the process is stopped. This setting is useful to ensure that the defragmentation process ends before the demand for server access is likely to increase.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger during the format. Otherwise defragmentation registers as a change by the Shadow Copy process. This increase in the number of changes forces Shadow Copy to delete snapshots as the limit for the cache file is reached.

CAUTION:

Allocation unit size cannot be altered without reformatting the drive. Data on a reformatted drive cannot be recovered.

For more information about disk defragmentation, read the online help.

Disk quotas

Disk quotas track and control disk space use in volumes.

NOTE:

To limit the size of a folder or share, see "Quota management" on page 56 .

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a

per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

NOTE:

When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

For more information about disk quotas, read the online help.

Adding storage

Expansion is the process of adding physical disks to an array that has already been configured. Extension is the process of adding new storage space to an existing logical drive on the same array, usually after the array has been expanded.

Storage growth may occur in three forms:

- Extend unallocated space from the original logical disks or LUNs.
- Alter LUNs to contain additional storage.
- Add new LUNs to the system.

The additional space is then extended through a variety of means, depending on which type of disk structure is in use.

NOTE:

This section addresses only single storage system node configurations. If your server has Windows Storage Server 2008 Enterprise Edition, see the Cluster Administration chapter for expanding and extending storage in a cluster environment.

Expanding storage

Expansion is the process of adding physical disks to an array that has already been configured. The logical drives (or volumes) that exist in the array before the expansion takes place are unchanged, because only the amount of free space in the array changes. The expansion process is entirely independent of the operating system.

NOTE:

See your storage array hardware user documentation for further details about expanding storage on the array.

Extending storage using Windows Storage Utilities

Volume extension grows the storage space of a logical drive. During this process, the administrator adds new storage space to an existing logical drive on the same array, usually after the array has been expanded. An administrator may have gained this new storage space by either expansion or by deleting another logical drive on the same array. Unlike drive expansion, the operating system must be aware of changes to the logical drive size.

You extend a volume to:

- Increase raw data storage
- Improve performance by increasing the number of spindles in a logical drive volume
- Change fault-tolerance (RAID) configurations

For more information about RAID levels, see the *Smart Array Controller User Guide*, or the document titled *Assessing RAID ADG vs. RAID 5 vs. RAID 1+0*. Both are available at the Smart Array controller web page or at http://h18000.www1.hp.com/products/servers/proliantstorage/arraycontrollers/documentation.html.

Extend volumes using Disk Management

The Disk Management snap-in provides management of hard disks, volumes or partitions. It can be used to extend a dynamic volume only.

NOTE:

Disk Management cannot be used to extend basic disk partitions.

Guidelines for extending a dynamic volume:

- Use the Disk Management utility.
- You can extend a volume only if it does not have a file system or if it is formatted NTFS.
- You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend striped volumes, mirrored volumes, or RAID 5 volumes.

For more information, see the Disk Management online help.

Expanding storage for EVA arrays using Command View EVA

Presenting a virtual disk offers its storage to a host. To make a virtual disk available to a host, you must present it. You can present a virtual disk to a host during or after virtual disk creation. The virtual disk must be completely created before the host presentation can occur. If you choose host presentation during virtual disk creation, the management agent cannot complete any other task until that virtual disk is created and presented. Therefore, HP recommends that you wait until a virtual disk is created before presenting it to a host.

For more information, see the *HP StorageWorks Command View EVA User Guide*.

Expanding storage using the Array Configuration Utility

The Array Configuration Utility enables online capacity expansion of the array and logical drive for specific MSA storage arrays, such as the MSA1000 and MSA1500. For more information, use the ACU online help, or the procedures to “Expand Array” in the *HP Array Configuration Utility User Guide*

Expand logical drive

This option in the ACU increases the storage capacity of a logical drive by adding unused space on an array to the logical drive on the same array. The unused space is obtained either by expanding an array or by deleting another logical drive on the same array. For more information, use the ACU online help, or the "Extend logical drive" procedure in the *HP Array Configuration Utility User Guide*

Volume shadow copies

NOTE:

Select storage systems can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses using shadow copies in a non-clustered environment.

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume.

A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the shadow copy mechanism is managed at the server, previous versions of files and folders are only available over the network from clients, and are seen on a per folder or file level, and not as an entire volume.

The shadow copy feature uses data blocks. As changes are made to the file system, the Shadow Copy Service copies the original blocks to a special cache file to maintain a consistent view of the file at a particular point in time. Because the snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space because blocks are not moved until an update to the disk occurs.

By using shadow copies, a storage system can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted. Previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file. A previous version of that file can be accessed.
- Compare several versions of a file while working. Use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. Because a snapshot only contains a portion of the original data blocks, shadow copies cannot protect against data loss due to media failures. However, the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow copy planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?

- How frequently will shadow copies be made?

Identifying the volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.

NOTE:

Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

Shadow copies are designed for volumes that store user data such as home directories and My Documents folders that are redirected by using Group Policy or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the `\\servername\sharename` path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.

NOTE:

Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating disk space

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily. If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, no shadow copy is created.

Administrators should also consider user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

NOTE:

Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 350 megabytes (MB). The default storage size is 10 percent of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the *storage* volume instead of the *source* volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.

△ CAUTION:

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

Identifying the storage area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on `H:\`, another volume such as `S:\`can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used storage systems.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

By keeping the shadow copy on the same volume, there is a potential gain in ease of setup and maintenance; however, there may be a reduction in performance and reliability.

△ CAUTION:

If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

Determining creation frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a maximum of 64 shadow copies per volume, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the storage system creates shadow copies at 0700 and 1200, Monday through Friday. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs.

Shadow copies and drive defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Using this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise, the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

NOTE:

To check the cluster size of a volume, use the `fsutil fsinfo ntfsinfo` command. To change the cluster size on a volume that contains data, back up the data on the volume, reformat it using the new cluster size, and then restore the data.

Mounted drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder `F:\data\users`, and the `Users` folder is a mount point for `G:\`. If shadow copies are enabled on both `F:\` and `G:\`, `F:\data` is shared as `\\server1\data`, and `G:\data\users` is shared as `\\server1\users`. In this example, users can access previous versions of `\\server1\data` and `\\server1\users` but not `\\server1\data\users`.

Managing shadow copies

The vssadmin tool provides a command line capability to create, list, resize, and delete volume shadow copies.

The system administrator can make shadow copies available to end users through a feature called "Shadow Copies for Shared Folders." The administrator uses the Properties menu (see Figure 5) to turn on the Shadow Copies feature, select the volumes to be copied, and determine the frequency with which shadow copies are made.

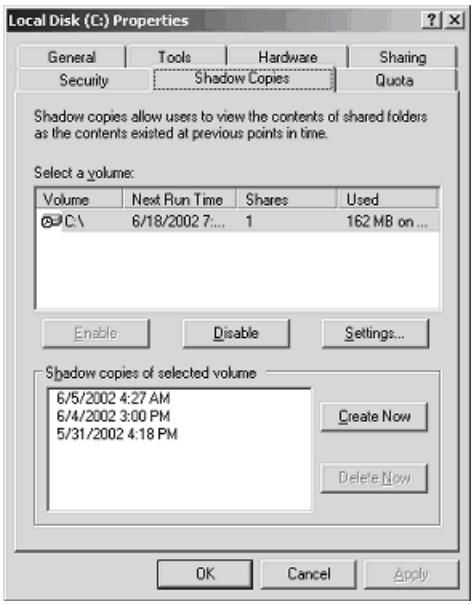

Figure 5 System administrator view of Shadow Copies for Shared Folders

The shadow copy cache file

The default shadow copy settings allocate 10 percent of the source volume being copied (with a minimum of 350 MB), and store the shadow copies on the same volume as the original volume. (See Figure 6). The cache file is located in a hidden protected directory titled "System Volume Information" off of the root of each volume for which shadow copy is enabled.

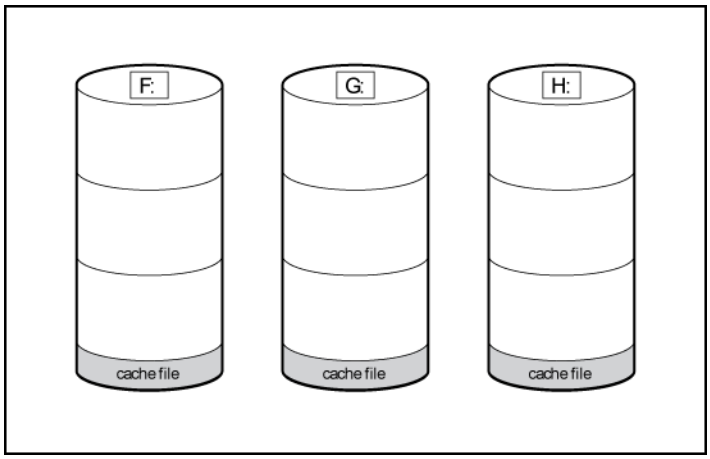

Figure 6 Shadow copies stored on a source volume

The cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. (See Figure 7).

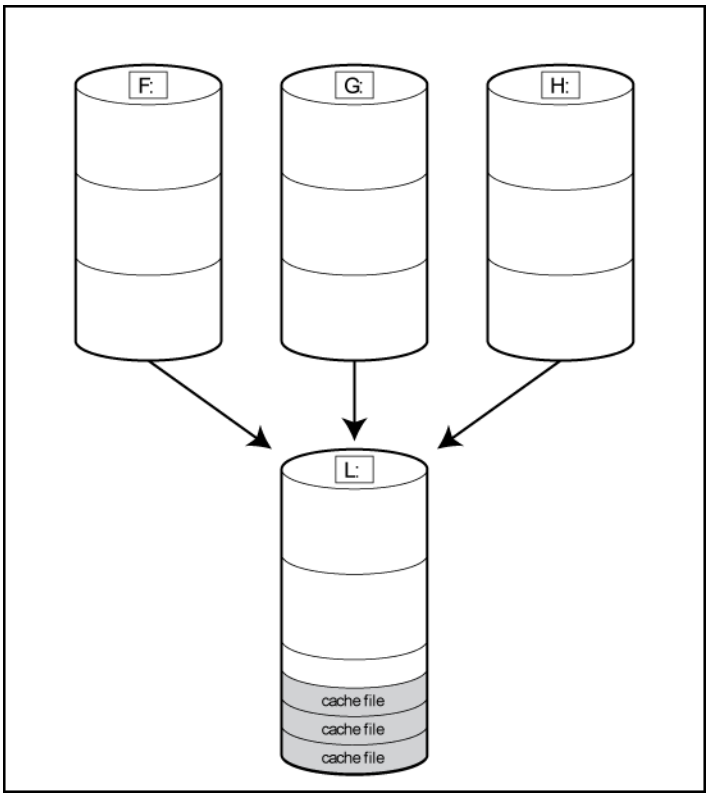

Figure 7 Shadow copies stored on a separate volume

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space, limits can generally be set higher, or set to No Limit. See the online help for instructions on altering the cache file location.

△ CAUTION:

If the data on the separate volume L: is lost, the shadow copies cannot be recovered.

Enabling and creating shadow copies

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume.
- Sets the maximum storage space for the shadow copies.
- Schedules shadow copies to be made at 7 a.m. and 12 noon on weekdays.

NOTE:

Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

NOTE:

After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See "Viewing shadow copy properties" on page 40.

Viewing a list of shadow copies

To view a list of shadow copies on a volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select **Shadow Copies** tab.

All shadow copies are listed, sorted by the date and time they were created.

NOTE:

It is also possible to create new shadow copies or delete shadow copies from this page.

Set schedules

Shadow copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow copy schedule to allow for these differences.

Do not schedule shadow copies more frequently than once per hour.

NOTE:

When deleting a shadow copy schedule, that action has no effect on existing shadow copies.

Viewing shadow copy properties

The Shadow Copy Properties page lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

NOTE:

For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. Managing the cache files on a separate disk is recommended.

CAUTION:

Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted.

Redirecting shadow copies to an alternate volume

IMPORTANT:

Shadow copies must be initially disabled on the volume before redirecting to an alternate volume. If shadow copies are enabled and you disable them, a message appears informing you that all existing shadow copies on the volume will be permanently deleted.

To redirect shadow copies to an alternate volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select the **Shadow Copies** tab.
5. Select the volume that you want to redirect shadow copies from and ensure that shadow copies are disabled on that volume; if enabled, click **Disable**.
6. Click **Settings**.
7. In the **Located on this volume** field, select an available alternate volume from the list.

NOTE:

To change the default shadow copy schedule settings, click **Schedule**.

8. Click **OK**.
9. On the **Shadow Copies** tab, ensure that the volume is selected, and then click **Enable**.

Shadow copies are now scheduled to be made on the alternate volume.

Disabling shadow copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

CAUTION:

When the Shadow Copies Service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

Managing shadow copies from the storage system desktop

To access shadow copies from the storage system desktop:

The storage system desktop can be accessed by using Remote Desktop to manage shadow copies.

1. On the storage system desktop, double-click **My Computer**.
2. Right-click the volume name, and select **Properties**.
3. Click the **Shadow Copies** tab. See Figure 8.

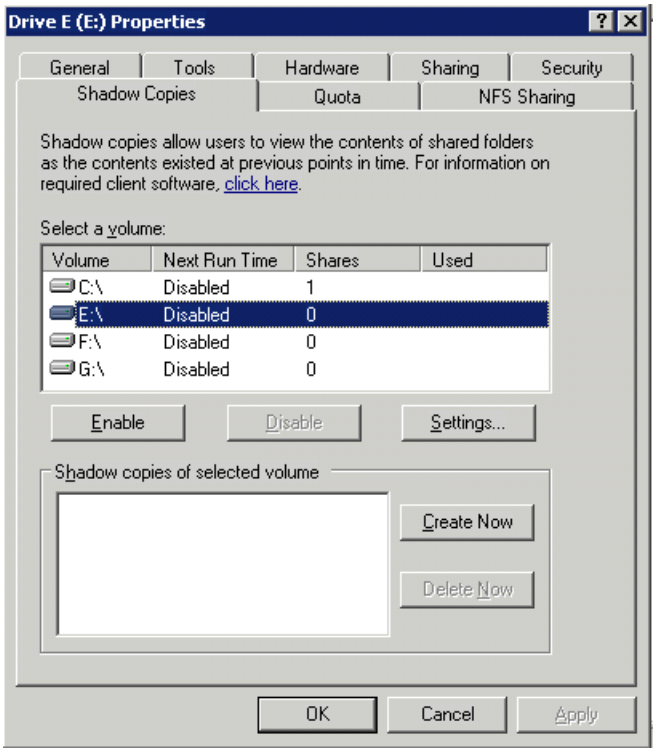

Figure 8 Accessing shadow copies from My Computer

Shadow Copies for Shared Folders

Shadow copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this includes HTTP, FTP, AppleTalk, and NetWare Shares. For SMB support, a client-side application denoted as Shadow Copies for Shared Folders is required. The client-side application is currently only available for Windows XP and Windows 2000 SP3+.

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.

NOTE:

Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.

NOTE:

Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files for these users.

SMB shadow copies

Windows users can independently access previous versions of files stored on SMB shares by using the Shadow Copies for Shared Folders client. After the Shadow Copies for Shared Folders client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its Properties window, clicking the **Previous Versions** tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

Shadow Copies for Shared Folders preserves the permissions set in the access control list (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

The Shadow Copies for Shared Folders client pack installs a **Previous Versions** tab in the **Properties** window of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting **View**, **Copy**, or **Restore** from the **Previous Versions** tab. (See Figure 9). Both individual files and folders can be restored.

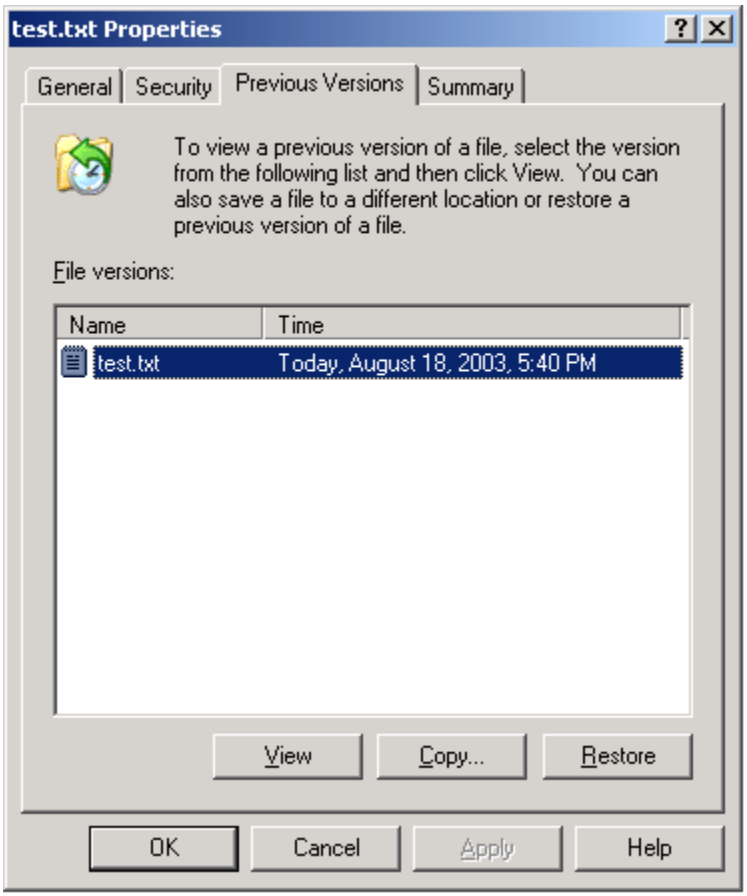

Figure 9 Client GUI

When users view a network folder hosted on the storage system for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history—a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS shadow copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format .@GMT-YYYY.MM.DD-HH:MM:SS. To prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named “NFSShare” with three shadow copies, taken on April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare

.@GMT-2003.04.27-04:00:00

.@GMT-2003.04.28-04:00:00

.@GMT-2003.04.29-04:00:00

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of files or folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation
- Accidental file replacement, which may occur if a user selects Save instead of Save As
- File corruption

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

Recovering a deleted file or folder

To recover a deleted file or folder within a folder:

1. Access to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file is selected.
3. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
4. Select the version of the folder that contains the file before it was deleted, and then click **View**.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Click **Restore** to restore the file or folder to its original location. Click **Copy...** to allow the placement of the file or folder to a new location.

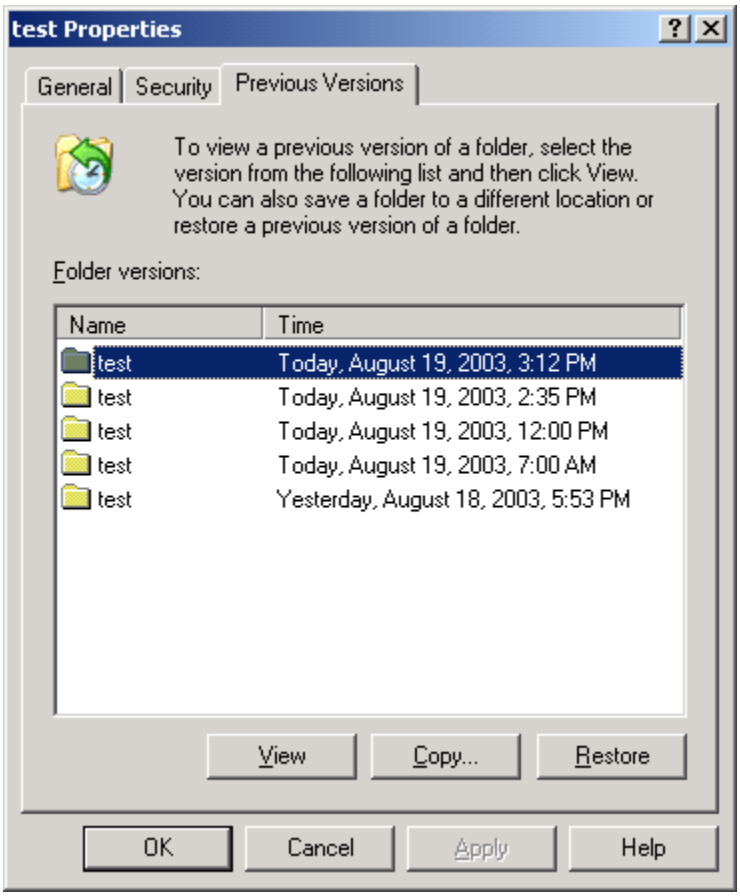

Figure 10 Recovering a deleted file or folder

Recovering an overwritten or corrupted file

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file:

1. Right-click the overwritten or corrupted file, and then click **Properties**.
2. Click **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

Recovering a folder

To recover a folder:

1. Position the cursor so that it is over a blank space in the folder to be recovered. If the cursor hovers over a file, that file is selected.
2. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
3. Click either **Copy...** or **Restore**.

   Clicking **Restore** enables the user to recover everything in that folder as well as all subfolders. Clicking **Restore** does not delete any files.

Backup and shadow copies

Shadow copies are only available on the network via the client application, and only at a file or folder level as opposed to the entire volume. Hence, the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, shadow copies are available for backup in two situations. If the backup software in question supports the use of shadow copies and can communicate with underlying block device, it is supported, and the previous version of the file system will be listed in the backup application as a complete file system snapshot. If the built-in backup application NTbackup is used, the backup software forces a snapshot, and then uses the snapshot as the means for backup. The user is unaware of this activity and it is not self-evident although it does address the issue of open files.

Shadow Copy Transport

Shadow Copy Transport provides the ability to transport data on a Storage Area Network (SAN). With a storage array and a VSS-aware hardware provider, it is possible to create a shadow copy on one server and import it on another server. This process, essentially "virtual" transport, is accomplished in a matter of minutes, regardless of the size of the data.

A shadow copy transport can be used for a number of purposes, including:

- Tape backups

  An alternative to traditional backup to tape processes is transport of shadow copies from the production server onto a backup server, where they can then be backed up to tape. Like the other two alternatives, this option removes backup traffic from the production server. While some backup applications might be designed with the hardware provider software that enables transport, others are not. The administrator should determine whether or not this functionality is included in the backup application.

- Data mining

  The data in use by a particular production server is often useful to different groups or departments within an organization. Rather than add additional traffic to the production server, a shadow copy of the data can be made available through transport to another server. The shadow copy can then be processed for different purposes, without any performance impact on the original server.

The transport process is accomplished through a series of DISKRAID command steps:

1. Create a shadow copy of the source data on the source server (read-only).
2. Mask off (hide) the shadow copy from the source server.
3. Unmask the shadow copy to a target server.
4. Optionally, clear the read-only flags on the shadow copy.

The data is now ready to use.

Folder and share management

The storage system supports several file-sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This section discusses overview information as well as procedures for the setup and management of the file shares for the supported protocols. Security at the file level and at the share level is also discussed.

NOTE:

Select servers can be deployed in a clustered or non-clustered configuration. This section discusses share setup for a non-clustered deployment.

Folder management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Folders can be managed using Server Manager. Tasks include:

- Accessing a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder

Managing file-level permissions

Security at the file level is managed using Windows Explorer.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, access the folder or file that needs to be changed, and then right-click the folder.

2. Click **Properties**, and then click the **Security** tab.

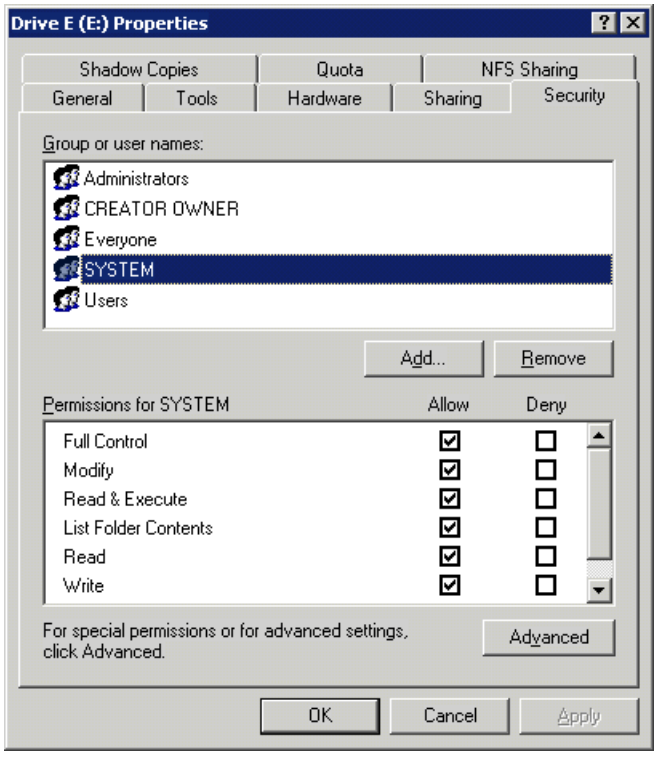

Figure 11 Properties dialog box, Security tab

Several options are available on the **Security** tab:

- To add users and groups to the permissions list, click **Add**. Follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group, and then click **Remove**.
- The center section of the **Security** tab lists permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file-access levels.

3. To modify ownership of files, or to modify individual file access level permissions, click **Advanced**.

   Figure 12 illustrates the properties available on the **Advanced Security Settings** dialog box.

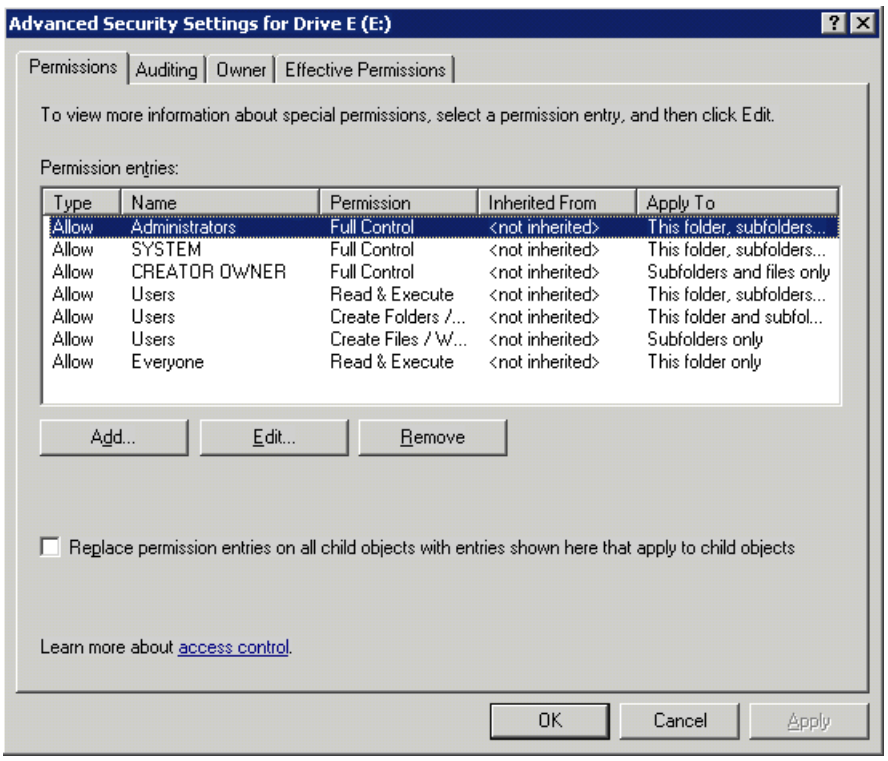

Figure 12 Advanced Security settings dialog box, Permissions tab

Other functionality available in the **Advanced Security Settings** dialog box is illustrated in Figure 12 and includes:

- Add a new user or group—Click **Add**, and then follow the dialog box instructions.
- Remove a user or group— Click **Remove**.
- Replace permission entries on all child objects with entries shown here that apply to child objects—This allows all child folders and files to inherit the current folder permissions by default.
- Modify specific permissions assigned to a particular user or group—Select the desired user or group, and then click **Edit**.

4. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. Figure 13 illustrates the **Edit** screen and some of the permissions.

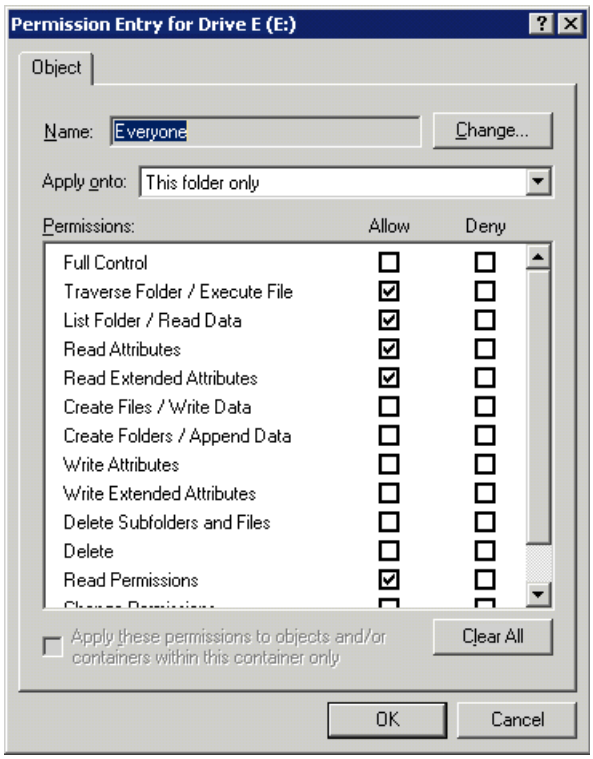

Figure 13 User or group Permission Entry dialog box

Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the **Advanced Security Settings Auditing** tab.

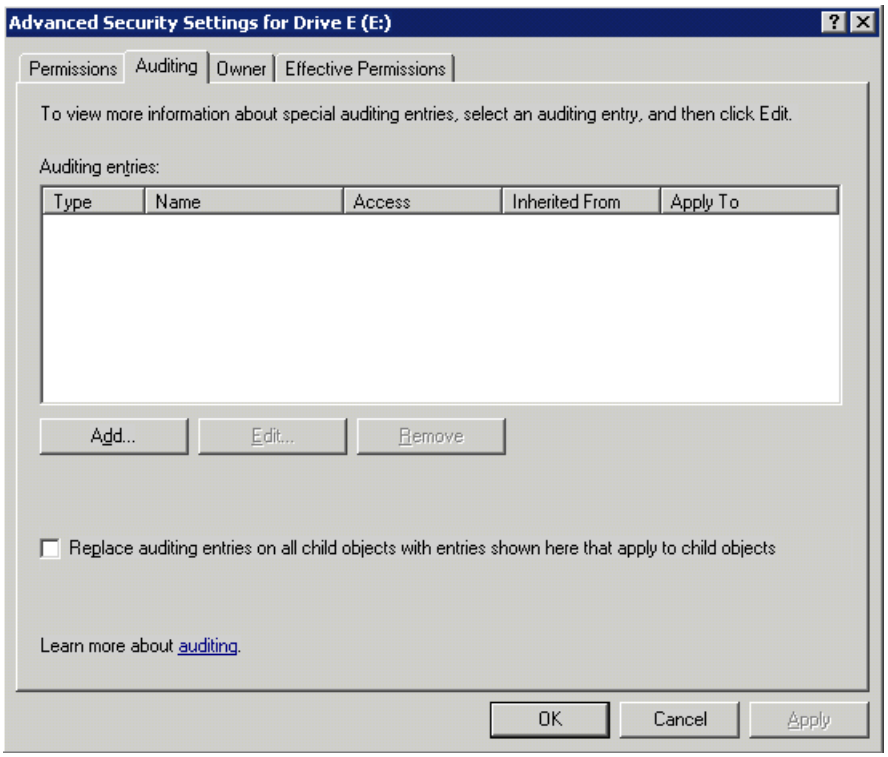

Figure 14 Advanced Security Settings dialog box, Auditing tab

5. Click **Add** to display the Select User or Group dialog box.

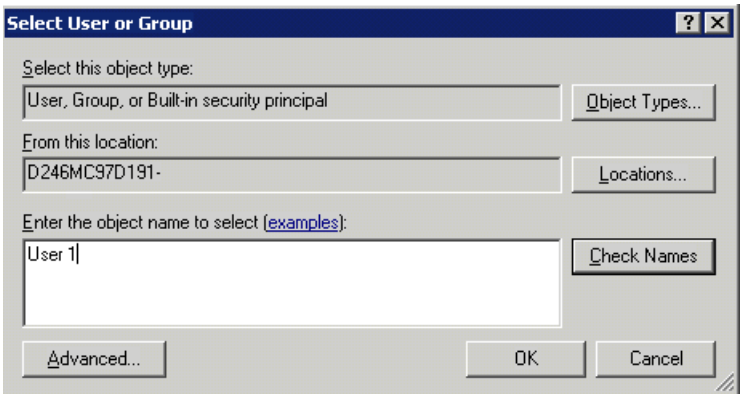

Figure 15 Select User or Group dialog box

NOTE:

Click Advanced to search for users or groups.

6. Select the user or group.

7. Click **OK**.

   The **Auditing Entry** dialog box is displayed.

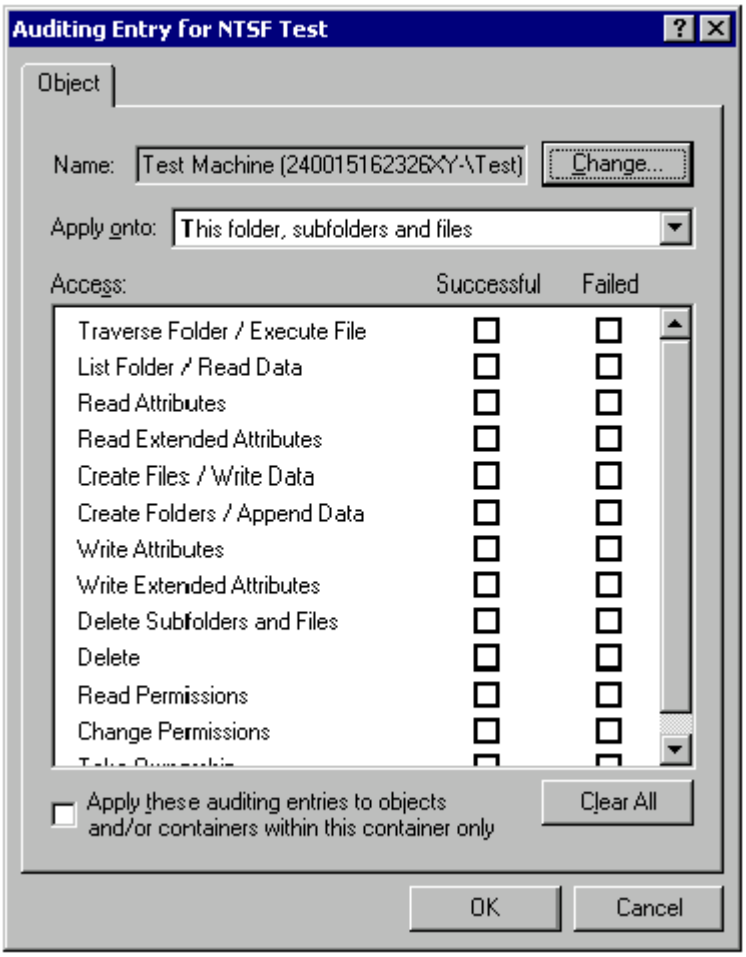

Figure 16 Auditing Entry dialog box for folder name NTFS Test

8. Select the desired **Successful** and **Failed** audits for the user or group.
9. Click **OK**.

NOTE:

Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the storage system.

The **Owner** tab allows taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files, and then manually apply the appropriate security configurations.

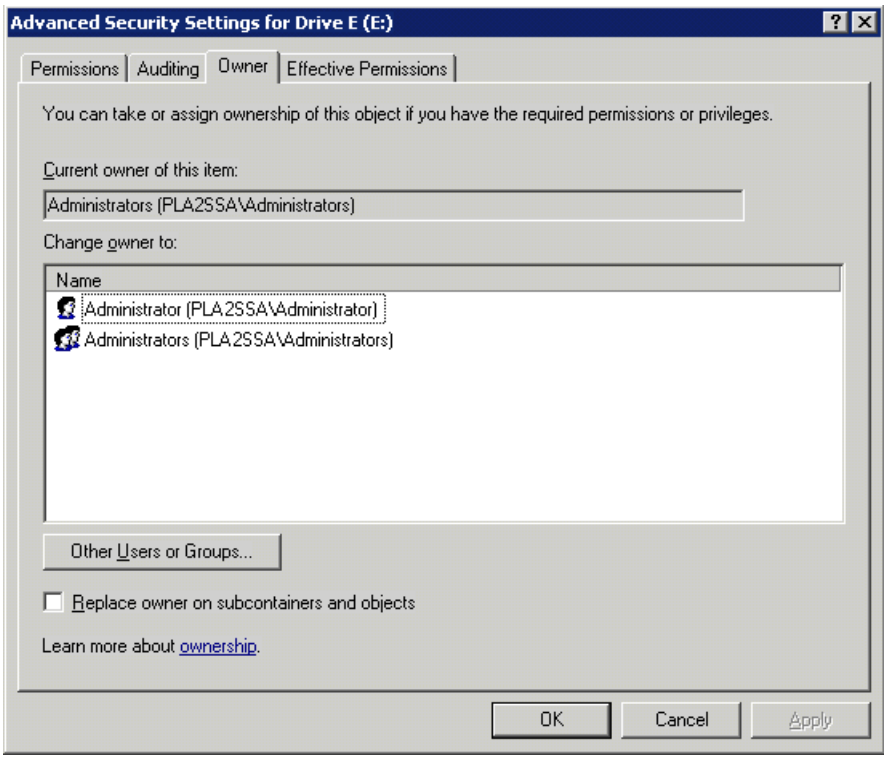

Figure 17 Advanced Security Settings dialog box, Owner tab

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Click the appropriate user or group in the **Change owner to** list.
2. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.
3. Click **OK**.

Share management

There are several ways to set up and manage shares. Methods include using Windows Explorer, a command line interface, or Server Manger.

NOTE:

Select servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment.

As previously mentioned, the file-sharing security model of the storage system is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security.

Share considerations

Planning the content, size, and distribution of shares on the storage system can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature, or of having very few shares of a generic nature. For example, shares for general use are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. For example, if it is sufficient to create a single share for user home directories, create a "homes" share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the storage system is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top-level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating local file system security into Windows domain environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the storage system can be given access permissions to shares managed by the device. The domain name of the storage system supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL, and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

NOTE:

Share permissions and file-level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file-level permissions override the share permissions.

Comparing administrative (hidden) and standard shares

CIFS supports both administrative shares and standard shares.

- Administrative shares are shares with a last character of $. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server.

- Standard shares are shares that do not end in a $ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The storage system supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the $ character when setting up the share. Do not type a $ character at the end of the share name when creating a standard share.

Managing shares

Shares can be managed using Server Manager. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties
- Publishing in DFS

NOTE:

These functions can operate in a cluster on select servers, but should only be used for non-cluster-aware shares. Use Cluster Administrator to manage shares for a cluster. The page will display cluster share resources.

CAUTION:

Before deleting a share, warn all users to exit that share and confirm that no one is using that share.

File Server Resource Manager

File Server Resource Manager (FSRM) is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. Some of the tasks you can perform are:

- Quota management
- File screening management
- Storage reports

Server Manager provides access to FSRM tasks.

For procedures and methods beyond what are described below, see the online help.

Quota management

On the Quota Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded.
- Generate auto quotas that apply to all existing folders in a volume or folder, as well as to any new subfolders created in the future.
- Define quota templates that can be easily applied to new volumes or folders and that can be used across an organization.

File screening management

On the File Screening Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create file screens to control the types of files that users can save and to send notifications when users attempt to save blocked files.
- Define file screening templates that can be easily applied to new volumes or folders and that can be used across an organization.
- Create file screening exceptions that extend the flexibility of the file screening rules.

Storage reports

On the Storage Reports node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Schedule periodic storage reports that allow you to identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or a selected group of users.
- Generate storage reports instantly.

Other Windows disk and data management tools

When you install certain tools, such as Windows Support Tools or Windows Resource Kit Tools, information about these tools might appear in Help and Support Center. To see the tools that are available to you, look in the Help and Support Center under **Support Tasks**, click **Tools**, and then click **Tools by Category**.

NOTE:

The Windows Support Tools and Windows Resource Kit Tools, including documentation for these tools, are available in English only. If you install them on a non-English language operating system or on an operating system with a Multilingual User Interface Pack (MUI), you see English content mixed with non-English content in Help and Support Center. To see the tools that are available to you, click **Start**, click **Help and Support Center**, and then, under **Support Tasks**, click **Tools**.

Additional information and references for file services

Backup

HP recommends that you back up the print server configuration whenever a new printer is added to the network and the print server configuration is modified.

HP StorageWorks Library and Tape Tools

HP StorageWorks Library and Tape Tools (L&TT) provides functionality for firmware downloads, verification of device operation, maintenance procedures, failure analysis, corrective service actions, and some utility functions. It also provides seamless integration with HP hardware support by generating and e-mailing support tickets that deliver a snapshot of the storage system.

For more information, and to download the utility, see the StorageWorks L&TT web site at http://h18006.www1.hp.com/products/storageworks/ltt.

Antivirus

The server should be secured by installing the appropriate antivirus software.anything

4 Cluster administration

HP StorageWorks X3000 Network Storage Systems support clustering; HP StorageWorks X1000 Network Storage Systems do not.

One important feature of HP StorageWorks X3000 Network Storage System models is that they can operate as a single node or as a cluster. This chapter discusses cluster installation and cluster management issues.

Cluster overview

Up to eight server nodes can be connected to each other and deployed as a no single point of failure (NSPOF) cluster. Utilizing a private network allows communication amongst themselves in order to track the state of each cluster node. Each node sends out periodic messages to the other nodes; these messages are called heartbeats. If a node stops sending heartbeats, the cluster service fails over any resources that the node owns to another node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat stops. The other nodes detect the lack of the heartbeat and another node takes over ownership of the Quorum disk and the cluster.

Clustering servers greatly enhances the availability of file serving by enabling file shares to fail over to additional storage systems if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

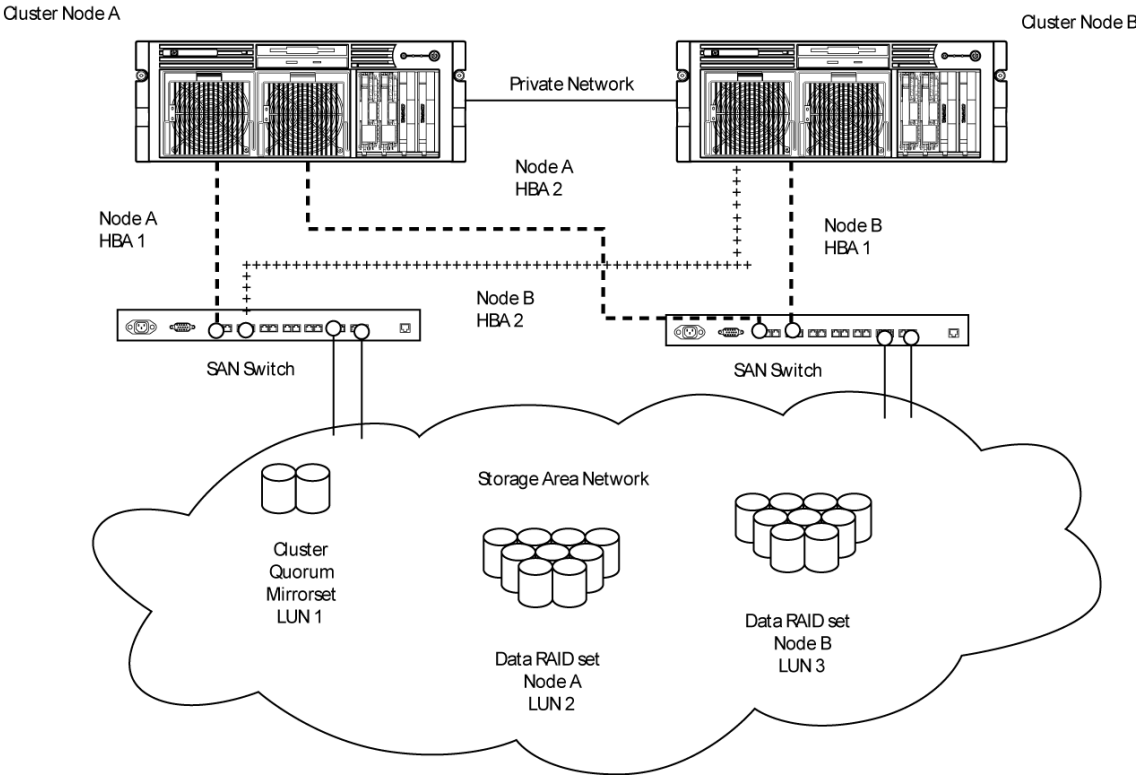

Figure 18 Storage system cluster diagram

Cluster terms and components

Nodes

The most basic parts of a cluster are the servers, referred to as nodes. A server node is any individual server in a cluster, or a member of the cluster.

Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Examples of cluster resources are IP addresses, network names, physical disk resources, and file shares. Resources represent individual system components. These resources are organized into groups and managed as a group. Some resources are created automatically by the system and other resources must be set up manually. Resource types include:

- IP address resource
- Cluster name resource
- Cluster quorum disk resource
- Physical disk resource
- Virtual server name resources

- CIFS file share resources
- NFS file share resources
- FTP file share resources
- iSCSI resources

Cluster groups

Cluster resources are placed together in cluster groups. Groups are the basic unit of failover between nodes. Resources do not fail over individually; they fail over with the group in which they are contained.

Virtual servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the storage systems can be distributed between the nodes of a cluster.

The creation of a virtual server allows resources dependent on the virtual server to fail over and fail back between the cluster nodes. Cluster resources are assigned to the virtual server to ensure non-disruptive service of the resources to the clients.

Failover and failback

Failover of cluster groups and resources happens:

- When a node hosting the group becomes inactive.
- When all of the resources within the group are dependent on one resource, and that resource fails.
- When an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owner's list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a failover occurs and the cluster is brought back to its original state, failback can occur automatically based on the policy. After a previously failed node comes online, the cluster service can fail back the groups to the original host. The failback policy must be set before the failover occurs so that failback works as intended.

Quorum disk

Each cluster must have a shared disk called the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by all nodes in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- Storing the most current version of the cluster database
- Guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster

Cluster concepts

Figure 19 illustrates a typical cluster configuration with the corresponding storage elements. The diagram progresses from the physical disks to the file shares, showing the relationship between both the cluster elements and the physical devices underlying them. While the diagram only illustrates two nodes, the same concepts apply for multi-node deployments.

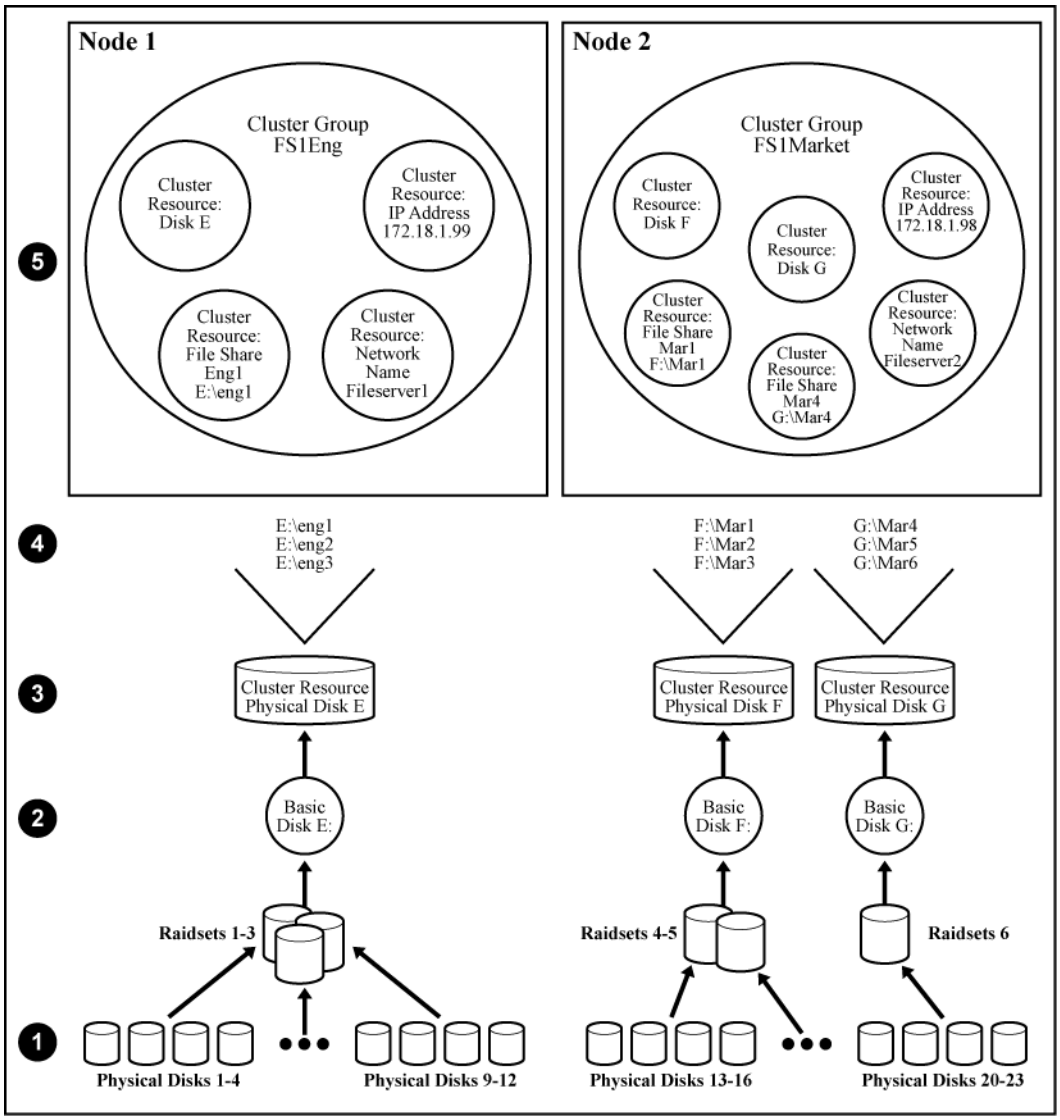

Figure 19 Cluster concepts diagram

Sequence of events for cluster resources

The sequence of events in the diagram includes:

1. Physical disks are combined into RAID arrays and LUNs.
2. LUNS are designated as basic disks, formatted, and assigned a drive letter via Disk Manager.

3. Physical Disk resources are created for each basic disk inside Cluster Administrator.
4. Directories and folders are created on assigned drives.
5. Cluster components (virtual servers, file shares) are created, organized in groups, and placed within the folders using Cluster Administrator exclusively.

Hierarchy of cluster resource components

Figure 19 depicts the cluster resource hierarchy as follows:

- Physical Disk resources are placed in a cluster group and relate to the basic disk. When a Physical Disk resource is created through Cluster Administrator, the resource should be inserted into an existing cluster group or a corresponding group should be created for the resource to reside in.
- File share resources are placed in a group and relate to the actual directory on the drive on which the share is being created.
- An IP Address resource is formed in the group and relates to the IP address by which the group's virtual server is identified on the network.
- A Network Name resource is formed in the group and relates to the name published on the network by which the group is identified.
- The Group is owned by one of the nodes of the cluster, but may transition to the other nodes during failover conditions.

The diagram illustrates a cluster containing two nodes. Each node has ownership of one group. Contained within each group are file shares that are known on the network by the associated Network Name and IP address. In the specific case of Node1, file share Eng1 relates to `E:\Eng1`. This file share is known on the network as `\\Fileserver1\Eng1` with an IP address of `172.18.1.99`.

For cluster resources to function properly, two very important requirements should be adhered to:

- Dependencies between resources of a group must be established. Dependencies determine the order of startup when a group comes online. In the above case, the following order should be maintained:
  1. File Share—Dependent on Physical Disk Resource and Network Name
  2. Network Name—Dependent on IP Address

  Failure to indicate the dependencies of a resource properly may result in the file share attempting to come online prior to the physical disk resource being available, resulting in a failed file share.
- Groups should have a Network Name resource and an IP Address resource. These resources are used by the network to give each group a virtual name. Without this virtual reference to the group, the only way to address a share that is created as a clustered resource is by node name. Physical node names do not transition during a failover, whereas virtual names do.

For example, if a client maps a network share to `\\Node1\Eng1` instead of `\\Fileserver1\Eng1`, when Node1 fails and Node2 assumes ownership, the map will become invalid because the reference in the map is to `\\Node1`. If the map were created to the virtual name and Node1 were to fail, the map would still exist when the group associated with Eng1 failed over to Node2.

The previous diagram is an example and is not intended to imply limitations of a single group or node. Groups can contain multiple physical disks resources and file shares and nodes can have multiple groups, as shown by the group owned by Node2.

Cluster planning

Requirements for taking advantage of clustering include:

- Storage planning
- Network planning
- Protocol planning

Storage planning

For clustering, a basic disk must be designated for the cluster and configured as the Quorum disk.

Additional basic disks are presented to each cluster node for data storage as physical disk resources. The physical disk resources are required for the basic disks to successfully work in a cluster environment, protecting it from simultaneous access from each node.

The basic disk must be added as a physical disk resource to an existing cluster group or a new cluster group needs to be created for the resource. Cluster groups can contain more than one physical disk resource depending on the site-specific requirements.

NOTE:

The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation or SAN zoning, or having only one node online at all times until the physical resource for the basic disk is established.

In preparing for the cluster installation:

- All shared disks, including the Quorum disk, must be accessible from all nodes. When testing connectivity between the nodes and the LUN, only one node should be given access to the LUN at a time.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Network planning

Clusters require more sophisticated networking arrangements than a stand alone storage system. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non domain environment.

All cluster deployments have at least six network addresses and four network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for virtual server A
- Cluster Interconnect static IP addresses for Node A and Node B

In multi-node deployments, additional network addresses are required. For each additional node, three static IP addresses are required.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the resources in the cluster group.

A cluster uses at least two network connections on each node:

- The private cluster interconnect or “heartbeat” crossover cable connects to one of the network ports on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The public client network subnet connects to the remaining network ports on each cluster node. The cluster node names and virtual server names have IP addresses residing on these subnets.

NOTE:

If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

Protocol planning

Not all file sharing protocols can take advantage of clustering. If a protocol does not support clustering, it will not have a cluster resource and will not failover with any cluster group. In the case of a failover, a client cannot use the virtual name or virtual IP address to access the share since the protocol cannot failover with the cluster group. The client must wait until the initial node is brought back online to access the share.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Table 4 Sharing protocol cluster support

Protocol	Client Variant	Cluster Aware (supports failover)	Supported on cluster nodes
CIFS/SMB	Windows NT Windows 2000 Windows 95 Windows 98 Windows ME	Yes	Yes
NFS	UNIX Linux	Yes	Yes
HTTP	Web	No	Yes
FTP	Many	Yes	Yes
NCP	Novell	No	Yes
AppleTalk	Apple	No	No
iSCSI	Standards-based iSCSI initiator	Yes	Yes

NOTE:

AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

Preparing for cluster installation

This section provides the steps necessary to cluster HP StorageWorks X3000 Network Storage Systems.

Before beginning installation

Confirm that the following specifications have been met before proceeding:

- The Quorum disk has been created from shared storage and is at least 50 MB. (500 MB is recommended.) Additional LUNs may also be presented for use as shared disk resources.
- Cluster configurations should be deployed with dual data paths for high availability. Dual data paths from each node enable a path failure to occur that does not force the failover of the node. Clusters can be configured with single path, but if a failure in the path does occur, all of the node resources will be failed to the non-affected node.

Using multipath data paths for high availability

HP recommends that cluster configurations be deployed with dual data paths for high availability. Clusters can be configured with single path, but if a failure in the path occurs, all of the node resources will be failed to the non-affected node. Pathing software is required in configurations where multipathing to the storage is desired or required. Multipathing software allows for datapath failure to occur without forcing a node failover.

Checklists for cluster server installation

These checklists assist in preparing for installation. Step-by-step instructions begin after the checklists.

Network requirements

- A unique NetBIOS cluster name
- For each node deployed in the cluster the following static IP addresses are required:
  - One for the network adapters on the private network
  - One for the network adapters on the public network
  - One for the virtual server itself

  A single static cluster IP address is required for the entire cluster.
- A domain user account for Cluster service (all nodes must be members of the same domain)
- Each node should have at least two network adapters—one for connection to the public network and the other for the node-to-node private cluster network. If only one network adapter is used for both connections, the configuration is unsupported. A separate private network adapter is required for HCL certification.

Shared disk requirements

NOTE:

Do not allow more than one node access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN zoning, or having only one node online at all times.

- All shared disks, including the Quorum disk, must be accessible from all nodes. When testing connectivity between the nodes and the LUN, only one node should be given access to the LUN at a time.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Cluster installation

During the installation process, nodes are shut down and rebooted. These steps guarantee that the data on disks that are attached to the shared storage bus is not lost or corrupted. This can happen when multiple nodes try to simultaneously write to the same disk that is not yet protected by the cluster software.

Use Table 5 to determine which nodes and storage devices should be presented during each step.

Table 5 Power sequencing for cluster installation

Step	Node 1	Additional Nodes	Storage	Comments
Setting up networks	On	On	Not Presented	Verify that all storage devices on the shared bus are not presented; Power on all nodes.
Setting up shared disks (including the Quorom disk)	On	Off	Presented	Shut down all nodes. Present the shared storage, then power on the first node.
Verifying disk configuration	Off	On	Presented	Shut down first node, power on next node. Repeat this process for all cluster nodes.
Configuring the first node	On	Off	Presented	Shut down all nodes; power on the first node.
Configuring additional nodes	On	On	Presented	Power on the next node after the first node is successfully configured. Complete this process for all cluster nodes.
Post-installation	On	On	Presented	At this point all cluster nodes should be on.

To configure the Cluster service on the storage system, an account must have administrative permissions on each node.

Setting up networks

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network.

Configuring the private network adapter

The following procedures are best practices provided by Microsoft and should be configured on the private network adapter.

- On the **General** tab of the private network adapter, ensure that only TCP/IP is selected.
- Ensure that the **Register this connection's address in DNS** is not selected in the DNS tab under advanced settings for Internet Protocol (TCP/IP) Properties.
- In all cases, set static IP addresses for the private network connector.

Configuring the public network adapter

While the public network adapter's IP address can be automatically obtained if a DHCP server is available, this is not recommended for cluster nodes. HP strongly recommends setting static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained though DHCP, access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost. Keep in mind that Cluster service recognizes only one network interface per subnet.

Renaming the local area connection icons

HP recommends changing the names of the network connections for clarity. The naming helps identify a network and correctly assign its role. For example, “Cluster interconnect” for the private network and “Public connection” for the public network.

Verifying connectivity and name resolution

To verify name resolution, ping each node from a client using the node's machine name instead of its IP address.

Verifying domain membership

All nodes in the cluster must be members of the same domain and able to access a domain controller and a DNS Server.

Setting up a cluster account

The Cluster service requires a domain user account under which the Cluster service can run. This user account must be created before installing Cluster service, because setup requires a user name and password. This user account should be a unique domain account created specifically to administer this cluster. This user account will need to be granted administrator privileges.

About the Quorum disk

HP makes the following Quorum disk recommendations:

- Dedicate a separate disk resource for a Quorum disk. Because the failure of the Quorum disk would cause the entire cluster to fail, HP strongly recommends that the disk resource be a RAID 1 configuration.
- Create a partition with a minimum of 50 megabytes (MB) to be used as a Quorum disk. HP recommends a Quorum disk be 500 MB.

HP recommends assigning the drive letter Q for the Quorum disk. It is also helpful to label the volume Quorum.

NOTE:

It is possible to change the Quorum disk by clicking the Quorum button. This displays a list of available disks that can be used for the Quorum disk. Select the appropriate disk, and then click **OK** to continue.

Configuring shared disks

Use the Windows Disk Management utility to configure additional shared disk resources. Verify that all shared disks are formatted as NTFS and are designated as Basic.

Additional shared disk resources are automatically added into the cluster as physical disk resources during the installation of cluster services.

Verifying disk access and functionality

Write a file to each shared disk resource to verify functionality.

At this time, shut down the first node, power on the next node and repeat the Verifying Disk Access and Functionality step above for all cluster nodes. When it has been verified that all nodes can read and write from the disks, turn off the cluster nodes and power on the first, and then continue with this guide.

Configuring cluster service software

Cluster Administrator (cluadmin) provides the ability to manage, monitor, create and modify clusters and cluster resources.

Using Cluster Administrator

Cluster Administrator shows information about the groups and resources on all of your clusters and specific information about the clusters themselves. A copy of Cluster Administrator is automatically installed on a cluster node when the Cluster service is installed.

Creating a cluster

During the creation of the cluster, Cluster Administrator will analyze and verify the hardware and software configuration and identify potential problems. A comprehensive and easy-to-read report is created, listing any potential configuration issues before the cluster is created.

Some issues that can occur are:

- No shared disk for the Quorum disk. A shared disk must be created with a NTFS partition at least 50 MB in size.

- Use of DHCP addresses for network connections. All Network adapters must be configured with static IP addresses in a cluster configuration.
- File Services for Macintosh and Service for NetWare are not supported in a cluster configuration.
- Dynamic Disks are not supported in a cluster configuration.
- Errors appear on a network adapter that is not configured or does not have an active link. If the network adapter is not going to be used it should be disabled.

Adding nodes to a cluster

Only the Quorum disk should be accessible by the new node while the new node is not a member of the cluster. The new node should not have access to the other LUNs in the cluster until after it has joined the cluster. After the node has joined the cluster, the LUNs may be presented to the new node. Move the physical disk resources over to the new node to confirm functionality.

△ CAUTION:

Presenting other LUNs to the non-clustered system could lead to data corruption.

Geographically dispersed clusters

Cluster nodes can be geographically dispersed to provide an additional layer of fault tolerance. Geographically dispersed clusters are also referred to as stretched clusters.

The following rules must be followed with geographically dispersed clusters:

- A network connection with latency of 500 milliseconds or less ensures that cluster consistency can be maintained. If the network latency is over 500 milliseconds, the cluster consistency cannot be easily maintained.
- All nodes must be on the same subnet.

Cluster groups and resources, including file shares

The Cluster Administrator tool provides complete online help for all cluster administration activities.

Cluster resources include administrative types of resources as well as file shares. The following paragraphs include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Cluster Administrator.

Cluster group overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.

△ CAUTION:

Do not delete or rename the Cluster Group or IP Address. Doing so results in losing the cluster and requires reinstallation of the cluster.

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the nodes.

Node-based cluster groups

Creating only one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node-based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

Load balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node with the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by Node A and the second cluster group owned by Node B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

File share resource planning issues

CIFS and NFS are cluster-aware protocols that support the Active/Active cluster model, allowing resources to be distributed and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for Node A and additional NFS file share resources can be assigned to a group owned by a virtual server for Node B.

Configuring the file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

Resource planning

1. Create a cluster group for each node in the cluster with an IP address resource and a network name resource.

   Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.

2. For NFS environments, configure the NFS server.

   NFS specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the Microsoft Services for NFS section within the "Other network file and print services" chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.

3. Create the file share resources.
4. Assign ownership of the file share resources to the resource groups.
   a. Divide ownership of the file share resource between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.
   b. Verify that the physical disk resource for this file share is also included in this group.
   c. Verify that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

Permissions and access rights on share resources

File Share and NFS Share permissions must be managed using the Cluster Administrator tool versus the individual shares on the file system themselves via Windows Explorer. Administering them through the Cluster Administrator tool allows the permissions to migrate from one node to other. In addition, permissions established using Explorer are lost after the share is failed or taken offline.

NFS cluster-specific issues

For convenience, all suggestions are listed below:

- Back up user and group mappings.

  To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- Map consistently.

  Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

- Map properly.
  - Valid UNIX users should be mapped to valid Windows users.
  - Valid UNIX groups should be mapped to valid Windows groups.
  - Mapped Windows user must have the "Access this computer from the Network privilege" or the mapping will be squashed.
  - The mapped Windows user must have an active password, or the mapping will be squashed.

- In a clustered deployment, create user name mappings using domain user accounts.

  Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.

- In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.

  If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.

- In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.

Example: If the password and group files are located at `c:\maps` on node 1, then they must also be at `c:\maps` on node 2. The contents of the password and group files must be the same on both nodes as well.

These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

Non cluster aware file sharing protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been saved to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fails over to the opposing node.

Adding new storage to a cluster

Present the new storage to one node in the cluster. This can be accomplished through selective storage presentation or through SAN zoning.

The tasks described below are used to add storage to a cluster. See the online help for clustering for additional details.

Creating physical disk resources

A physical disk resource must reside within a cluster group. An existing cluster group can be used or a new cluster group must be created. For information on creating disk resources, see the cluster online help topic *Physical Disk resource type*.

NOTE:

- Physical disk resources usually do not have any dependencies set.
- In multi-node clusters it is necessary to specify the node to move the group to. When a cluster group is moved to another node, all resources in that group are moved.
- When a physical disk resource is owned by a node, the disk appears as an unknown, unreadable disk to all other cluster nodes. This is a normal condition. When the physical disk resource moves to another node, the disk resource then becomes readable.

Creating file share resources

To create a file share resource, see two clustering online help topics:

- Create a cluster-managed file share
- Using a server cluster with large numbers of file shares

NOTE:

- A file share resource must reside in the same cluster group as the physical disk resource it will reside on.
- The physical disk resource specified in this step must reside in the same cluster group as specified in the beginning of this wizard.

Creating NFS share resources

To create an NFS share resource, see "MSNFS administration on a server cluster" on page 74.

Shadow copies in a cluster

It is recommended that the location of the cache file be placed on a separate disk from the original data. In this case, a physical disk resource for the cache file disk should be created in the same cluster group as the intended Shadow Copy resource and the volume for which snapshots will be enabled. The resource should be created prior to the establishment of Shadow Copies. The Shadow Copy resource should be dependent on both the original physical disk resource and the physical disk resource that contains the cache file.

For more information, see the following topics in the clustering online help:

- Using Shadow Copies of Shared Folders in a server cluster
- Enable Shadow Copies for shared folders in a cluster

Extend a LUN in a cluster

To extend a LUN on a storage array in a cluster, review the requirements and procedures from the storage array hardware provider for expanding or extending storage.

For additional information associated with extending a LUN in a cluster, see the following Microsoft Knowledge Base articles:

- How to extend the partition of a cluster shared disk
  http://support.microsoft.com/default.aspx?scid=kb;en-us;304736
- How to replace a disk that is in a cluster and use of the Cluster Recovery utility
  http://support.microsoft.com/kb/305793

MSNFS administration on a server cluster

The Microsoft Services for Network File System (NFS) online help provides server cluster information for the following topics:

- Configuring shared folders on a server cluster
  - Configuring an NFS share as a cluster resource
  - Modifying an NFS shared cluster resource
  - Deleting an NFS shared cluster resource
- Using Microsoft Services for NFS with server clusters
  - Understanding how Server for NFS works with server clusters
  - Using Server for NFS on a server cluster

- Configuring User Name Mapping on a server cluster

For further details, see the online help for Microsoft Services for Network File System.

Best practices for running Server for NFS in a server cluster

- Stop Server for NFS before stopping the server cluster.
- Ensure share availability when a node fails.
- Use the appropriate tool to manage Network File System (NFS) share cluster resources.
- Avoid conflicting share names.
- Ensure the availability of audit logs.
- Move file shares or take them offline before stopping Server for NFS.
- Take resources offline before modifying.
- Administer Server for NFS only from computers in a trusted domain.
- Restart the Server for NFS service after the cluster service restarts.
- Choose the appropriate sharing mode.
- Use the command line properly when creating or modifying NFS share cluster resources.
- Use hard mounts.
- Use the correct virtual server name.

Print services in a cluster

The Windows Storage Server 2008 Cluster service implementation increases availability of critical print servers. A print spooler service on a clustered print server may be hosted on any of the nodes in the cluster. As with all cluster resources, clients should access the print server by its virtual network name or virtual IP address.

Creating a cluster printer spooler

Printer spoolers should be created in a separate group dedicated to this purpose for ease of management. For each printer spooler, a physical resource is required to instantiate the print spooler resource. In some cases, dedicated physical resources are not available and hence sharing of the physical resource among other members of the group is acceptable, remembering that all members of a group are managed as a unit. Hence, the group will failover and failback as a group.

To create a printer spooler:

1. Create a dedicated group (if desired).
2. Create a physical resource (disk) (if required, see note).
3. Create an IP address resource for the Virtual Server to be created (if required, see note).
4. Create a Virtual Server Resource (Network Name) (if required, see note).

NOTE:

If the printer spool resource is added to an existing group with a physical resource, IP address, and virtual server resource, steps 1-4 are not required.

5. Create a Print Spool resource.

6. To add a printer to the virtual server:
   a. Double-click the printers and faxes icon.
   b. Right-click the new screen, and then click **add printer**. A wizard starts.
   c. Click **create a new port**, and then click **Next**.
   d. Enter the IP address of the network printer.
   e. Update the Port Name if desired, click **Next**, and then click **Finish**.
   f. Select the appropriate driver, and then click **Next**.
   g. If presented with a dialog to replace the driver present, click **keep the driver**, and then click **Next**.
   h. Name the printer, and then click **Next**.
   i. Provide a share name for the printer for network access, and then click **Next**.
   j. Provide location information and comments, and then click **Next**.
   k. Click **Yes** to print a test page, click **Next**, and then click **Finish**.
   l. A dialog box appears regarding the test page. Select the appropriate answer.

The Printer Spool is now a clustered resource.

Advanced cluster administration procedures

Failing over and failing back

As previously mentioned, when a node goes offline, all resources dependent on that node are automatically failed over to another node. Processing continues, but in a reduced manner, because all operations must be processed on the remaining node(s). In clusters containing more than two nodes, additional fail over rules can be applied. For instance, groups can be configured to fail over different nodes to balance the additional work load imposed by the failed node. Nodes can be excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly the preferred owners list can be ordered, to provide an ordered list of failover nodes. Using these tools, the failover of resources can be controlled with in a multinode cluster to provide a controlled balanced failover methodology that balances the increased work load.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually.

NOTE:

If the storage system is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs.

Restarting one cluster node

△ CAUTION:

Restarting a cluster node should be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being restarted. Attached connections can be viewed through Server Manager on the storage system Desktop using Terminal Services. From Server Manager, select **File Sharing > Shared Folders > Sessions**.

The physical process of restarting one of the nodes of a cluster is the same as restarting a storage system in single node environment. However, additional caution is needed.

Restarting a cluster node causes all cluster resources served by that node to fail over to the other nodes in the cluster based on the failover policy in place. Until the failover process completes, any currently executing read and write operations will fail. Other node(s) in the cluster will be placed under a heavier load by the extra work until the restarted node comes up and the resources are moved back.

Shutting down one cluster node

△ CAUTION:

Shutting down a cluster node must be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being shutdown.

Shutting down a cluster node causes all cluster resources served by that node to fail over to the other node(s). This causes any currently executing client read and write operations to fail until the cluster failover process completes. The other node(s) are placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

Powering down the cluster

The power down process for the storage system cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.

△ CAUTION:

Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "Shutting down one cluster node." Only one cluster node should be shut down at a time.

Powering up the cluster

The power up process for the storage system cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.

△ CAUTION:

Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

Nodes should be powered up separately allowing one node to form the cluster prior to powering up the additional node(s). To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node. Wait for the node to come completely up before powering up the subsequent node(s).

   If more than one node is powered up at the same time, the first node that completes the sequence gains ownership of the cluster quorum and controls the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up additional cluster node(s).
2. Power up the additional cluster node(s). Each node should be allowed to start fully, prior to starting a subsequent node.

5 Troubleshooting, servicing, and maintenance

Troubleshooting the storage system

The "Support and troubleshooting" task at the HP Support & Drivers web site (http://www.hp.com/go/support) can be used to troubleshoot problems with the storage system. After entering the storage system name and designation (for example, ML110 G5 storage system) or component information (for example, Array Configuration Utility), use the following links for troubleshooting information:

- Download drivers and software—This area provides drivers and software for your operating system.
- Troubleshoot a problem—This area provides a listing of customer notices, advisories, and bulletins applicable for the product or component.
- Manuals—This area provides the latest user documentation applicable to the product or component. User guides can be a useful source for troubleshooting information. For most storage system hardware platforms, the following ProLiant server manuals may be useful for troubleshooting assistance:
  - **HP ProLiant Server User Guide or HP ProLiant Server Maintenance and Service Guide**.
    These guides contain specific troubleshooting information for the server.
  - **HP ProLiant Servers Troubleshooting Guide**
    The guide provides common procedures and solutions for many levels of troubleshooting with a ProLiant server. The guide is available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300504/c00300504.pdf.

IMPORTANT:

Some troubleshooting procedures found in ProLiant server guides may not apply to the HP StorageWorks X1000 and X3000 Network Storage Systems. If necessary, check with your HP Support representative for further assistance.

For software related components and issues, online help or user guide documentation may offer troubleshooting assistance. The release notes for the storage system product line is updated frequently. The document contains issues and workarounds to a number of categories for the storage systems.

Known issues and workarounds for the storage system products and the service release are addressed in release notes. To view the latest release notes, go to http://www.hp.com/go/nas, select your product family, product model, click **Support for your product**, and then click **Manuals**.

WEBES (Web Based Enterprise Services)

WEBES is a tool suite aimed at preventing or reducing your system's down time. The tool suite has the following components:

- CCAT (Computer Crash Analysis Tool)
- SEA (System Event Analyzer)

If you have a warranty or service contract with HP you are entitled to these tools free of charge. You must, however, upgrade the tools at least once a year because the software expires after one year. For more information about WEBES, see http://h18023.www1.hp.com/support/svctools/webes/.

To install WEBES on your storage system, run the setup executable located in the `C:\hpnas\Components\WEBES` folder.

Maintenance and service

HP provides specific documentation for maintaining and servicing your storage system and offers a customer self repair program.

Maintenance updates

Regular updates to the storage system are supplied on the HP StorageWorks Service Release DVD. The Service Release DVD can be obtained at http://www.software.hp.com.

Individual updates for each product are available for download from the HP Support web site at http://h18023.www1.hp.com/support/selfrepair/na/replace_part.asp.

System updates

System updates to the hardware (BIOS, firmware, drivers), critical updates, and hotfixes for the operating system and other related software updates are bundled on the Service Release DVD.

Firmware updates

Firmware is software that is stored in Read-Only Memory (ROM). Firmware is responsible for the behavior of the system when it is first switched on and for passing control of the server to the operating system. When referring to the firmware on the system board of the server, it is called the System ROM or the BIOS. When referring to the firmware on another piece of hardware configured in the server, it is called Option ROM. Storage systems have hard drives, Smart Array Controllers, Remote Insight Lights-Out Edition (RILOE), Remote Insight Lights-Out Edition II (RILOE II) and Integrated Lights-Out options that have firmware that can be updated.

It is important to update the firmware (also called "flashing the ROM") as part of regular server maintenance. In addition, checking for specific firmware updates in between regular updates helps to keep the server performing optimally. HP recommends checking for a firmware update before sending a part back to HP for replacement.

Certificate of Authenticity

The Certificate of Authenticity (COA) label is used to:

- Upgrade the factory-installed operating system using the Microsoft Upgrade program for license validation.
- Reinstall the operating system because of a failure that has permanently disabled it.

The COA label location varies by server model. On rack-mounted server models, the COA label is located either on the front section of the right panel or on the right front corner of the top panel. On tower models, the COA label is located toward the rear of the top panel of the server.

6 System recovery

This chapter describes how to use the System Recovery DVD that is provided with your storage system.

The System Recovery DVD

The HP StorageWorks Storage System Recovery DVD that is provided with your storage system allows you to install an image or recover from a catastrophic failure.

At any later time, you may boot from the DVD and restore the server to the factory condition. This allows you to recover the system if all other means to boot the server fail.

While the recovery process makes every attempt to preserve the existing data volumes, you should have a backup of your data if at all possible before recovering the system.

To restore a factory image

1. Insert the System Recovery DVD. The main window appears.
2. Choose **Restore Factory Image**.

NOTE:

If your storage system's DVD drive is not functional, or if your storage system does not have a DVD drive, you can run the system recovery from a bootable USB flash drive.

The DON'T ERASE partition

The DON'T ERASE logical disk supports the restoration process only and does not host a secondary operating system. Be sure to back up your user data, and then use the Recovery and Installation DVD to restore the server to the factory state.

Managing disks after a restoration

After a system has been restored, drive letters may be assigned to the wrong volume. Windows Storage Server 2008 assigns drive letters after the restoration in the order of discovery. To help maintain drive letter information, placing the drive letter into a volume label is recommended. To change the drive letters to the appropriate one, go into Disk Management and perform the following steps for each volume:

1. Right-click the volume that needs to be changed.
2. Select **Change drive Letter and Paths**.
3. In the **Change drive Letter and Paths** dialog box, select **Change**.

4. Select the appropriate drive letter, then click **OK**.
5. Click **Yes** to confirm the drive letter change.
6. Click **Yes** to continue. If the old drive letter needs to be re-used, reboot the server after clicking Yes.

A Regulatory compliance notices

This section contains regulatory notices for the HP StorageWorks family of products.

Regulatory compliance identification numbers

For the purpose of regulatory compliance certifications and identification, this product has been assigned a unique regulatory model number. The regulatory model number can be found on the product nameplate label, along with all required approval markings and information. When requesting compliance information for this product, always refer to this regulatory model number. The regulatory model number is not the marketing name or model number of the product.

Product specific information:

HP _____

Regulatory model number: _____

FCC and CISPR classification: _____

These products contain laser components. See Class 1 laser statement in the Laser compliance notices section.

Federal Communications Commission notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

FCC rating label

The FCC rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or ID on the label. Class A devices do not have an FCC logo or ID on the label. After you determine the class of the device, refer to the corresponding statement.

Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation

of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit that is different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Declaration of Conformity for products marked with the FCC logo, United States only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding this FCC declaration, contact us by mail or telephone:

- Hewlett-Packard Company P.O. Box 692000, Mail Stop 510101 Houston, Texas 77269-2000
- Or call 1-281-514-3333

Modification

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

When provided, connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Canadian notice (Avis Canadien)

Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union notice

Products bearing the CE marking comply with the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (in parentheses are the equivalent international standards and regulations):

- EN 55022 (CISPR 22)—Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11)—Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2)—Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3)—Power Line Flicker
- EN 60950 (IEC60950)—Product Safety

Japanese notices

Japanese VCCI-A notice

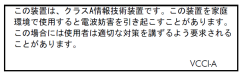

Japanese VCCI-B notice

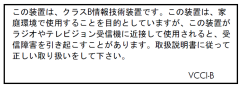

Japanese power cord statement

製品には、同梱された電源コードをお使い下さい。
同梱された電源コードは、他の製品では使用出来ません。

Please use the attached power cord.
The attached power cord is not allowed to use with other product.

Korean notices

Class A equipment

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Class B equipment

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

Taiwanese notices

BSMI Class A notice

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Taiwan battery recycle statement

Recovery mark: • Four-in-one recycling symbol	Recovery text: • "Please recycle waste batteries"
	廢電池請回收

Laser compliance notices

English laser notice

This device may contain a laser that is classified as a Class 1 Laser Product in accordance with U.S. FDA regulations and the IEC 60825-1. The product does not emit hazardous laser radiation.

⚠ WARNING!

Use of controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module enclosure. There are no user-serviceable components inside.
- Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.
- Allow only HP Authorized Service technicians to repair the unit.

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

Dutch laser notice

WAARSCHUWING: dit apparaat bevat mogelijk een laser die is geclassificeerd als een laserproduct van Klasse 1 overeenkomstig de bepalingen van de Amerikaanse FDA en de richtlijn IEC 60825-1. Dit product geeft geen gevaarlijke laserstraling af.

Als u bedieningselementen gebruikt, instellingen aanpast of procedures uitvoert op een andere manier dan in deze publicatie of in de installatiehandleiding van het laserproduct wordt aangegeven, loopt u het risico te worden blootgesteld aan gevaarlijke straling. Het risico van blootstelling aan gevaarlijke straling beperkt u als volgt:

- Probeer de behuizing van de module niet te openen. U mag zelf geen onderdelen repareren.
- Gebruik voor de laserapparatuur geen andere knoppen of instellingen en voer geen andere aanpassingen of procedures uit dan die in deze handleiding worden beschreven.
- Alleen door HP geautoriseerde technici mogen het apparaat repareren.

French laser notice

AVERTISSEMENT : cet appareil peut être équipé d'un laser classé en tant que Produit laser de classe 1 et conforme à la réglementation de la FDA américaine et à la norme 60825-1 de l'IEC. Ce produit n'émet pas de rayonnement dangereux.

L'utilisation de commandes, de réglages ou de procédures autres que ceux qui sont indiqués ici ou dans le manuel d'installation du produit laser peut exposer l'utilisateur à des rayonnements dangereux. Pour réduire le risque d'exposition à des rayonnements dangereux :

- Ne tentez pas d'ouvrir le boîtier renfermant l'appareil laser. Il ne contient aucune pièce dont la maintenance puisse être effectuée par l'utilisateur.
- Tout contrôle, réglage ou procédure autre que ceux décrits dans ce chapitre ne doivent pas être effectués par l'utilisateur.
- Seuls les Mainteneurs Agréés HP sont habilités à réparer l'appareil laser.

German laser notice

VORSICHT: Dieses Gerät enthält möglicherweise einen Laser, der nach den US-amerikanischen FDA-Bestimmungen und nach IEC 60825-1 als Laserprodukt der Klasse 1 zertifiziert ist. Gesundheitsschädliche Laserstrahlen werden nicht emittiert.

Die Anleitungen in diesem Dokument müssen befolgt werden. Bei Einstellungen oder Durchführung sonstiger Verfahren, die über die Anleitungen in diesem Dokument bzw. im Installationshandbuch des Lasergeräts hinausgehen, kann es zum Austritt gefährlicher Strahlung kommen. Zur Vermeidung der Freisetzung gefährlicher Strahlungen sind die folgenden Punkte zu beachten:

- Versuchen Sie nicht, die Abdeckung des Lasermoduls zu öffnen. Im Inneren befinden sich keine Komponenten, die vom Benutzer gewartet werden können.
- Benutzen Sie das Lasergerät ausschließlich gemäß den Anleitungen und Hinweisen in diesem Dokument.
- Lassen Sie das Gerät nur von einem HP Servicepartner reparieren.

Italian laser notice

AVVERTENZA: AVVERTENZA Questo dispositivo può contenere un laser classificato come prodotto laser di Classe 1 in conformità alle normative US FDA e IEC 60825-1. Questo prodotto non emette radiazioni laser pericolose.

L'eventuale esecuzione di comandi, regolazioni o procedure difformi a quanto specificato nella presente documentazione o nella guida di installazione del prodotto può causare l'esposizione a radiazioni nocive. Per ridurre i rischi di esposizione a radiazioni pericolose, attenersi alle seguenti precauzioni:

- Non cercare di aprire il contenitore del modulo. All'interno non vi sono componenti soggetti a manutenzione da parte dell'utente.
- Non eseguire operazioni di controllo, regolazione o di altro genere su un dispositivo laser ad eccezione di quelle specificate da queste istruzioni.
- Affidare gli interventi di riparazione dell'unità esclusivamente ai tecnici dell'Assistenza autorizzata HP.

Japanese laser notice

警告:本製品には、US FDA規則およびIEC 60825-1に基づくClass 1レーザー製品が含まれている場合があります。本製品は人体に危険なレーザー光は発しません。

本書およびレーザー製品のインストール ガイドに示されている以外の方法で制御、調整、使用した場合、人体に危険な光線にさらされる場合があります。人体に危険な光線にさらされないため、以下の項目を守ってください。

- モジュール エンクロージャを開けないでください。ユーザーが取り扱えるコンポーネントは含まれていません。
- 本書に示されている以外の方法で、レーザー デバイスを制御、調整、使用しないでください。
- HPの正規サービス技術者のみが本ユニットの修理を許可されています。

Spanish laser notice

ADVERTENCIA: Este dispositivo podría contener un láser clasificado como producto de láser de Clase 1 de acuerdo con la normativa de la FDA de EE.UU. e IEC 60825-1. El producto no emite radiaciones láser peligrosas.

El uso de controles, ajustes o manipulaciones distintos de los especificados aquí o en la guía de instalación del producto de láser puede producir una exposición peligrosa a las radiaciones. Para evitar el riesgo de exposición a radiaciones peligrosas:

- No intente abrir la cubierta del módulo. Dentro no hay componentes que el usuario pueda reparar.
- No realice más operaciones de control, ajustes o manipulaciones en el dispositivo láser que los aquí especificados.
- Sólo permita reparar la unidad a los agentes del servicio técnico autorizado HP.

Recycling notices

English notice

Disposal of Waste Equipment by Users in Private Households in the European Union

This symbol means do not dispose of your product with your other household waste. Instead, you should protect human health and the environment by handing over your waste equipment to a designated collection point for the recycling of waste electrical and electronic equipment. For more information, please contact your household waste disposal service

Bulgarian notice

Изхвърляне на отпадъчно оборудване от потребители в частни домакинства в Европейския съюз
Този символ върху продукта или опаковката му показва, че продуктът не трябва да се изхвърля заедно с другите битови отпадъци. Вместо това, трябва да предпазите човешкото здраве и околната среда, като предадете отпадъчното оборудване в предназначен за събирането му пункт за рециклиране на неизползваемо електрическо и електронно борудване. За допълнителна информация се свържете с фирмата по чистота, чиито услуги използвате.

Czech notice

Likvidace zařízení v domácnostech v Evropské unii
Tento symbol znamená, že nesmíte tento produkt likvidovat spolu s jiným domovním odpadem. Místo toho byste měli chránit lidské zdraví a životní prostředí tím, že jej předáte na k tomu určené sběrné pracoviště, kde se zabývají recyklací elektrického a elektronického vybavení. Pro více informací kontaktujte společnost zabývající se sběrem a svozem domovního odpadu.

Danish notice

Bortskaffelse af brugt udstyr hos brugere i private hjem i EU
Dette symbol betyder, at produktet ikke må bortskaffes sammen med andet husholdningsaffald. Du skal i stedet den menneskelige sundhed og miljøet ved at afl evere dit brugte udstyr på et dertil beregnet indsamlingssted for af brugt, elektrisk og elektronisk udstyr. Kontakt nærmeste renovationsafdeling for yderligere oplysninger.

Dutch notice

Inzameling van afgedankte apparatuur van particuliere huishoudens in de Europese Unie
Dit symbool betekent dat het product niet mag worden gedeponeerd bij het overige huishoudelijke afval. Bescherm de gezondheid en het milieu door afgedankte apparatuur in te leveren bij een hiervoor bestemd inzamelpunt voor recycling van afgedankte elektrische en elektronische apparatuur. Neem voor meer informatie contact op met uw gemeentereinigingsdienst.

Estonian notice

Äravisatavate seadmete likvideerimine Euroopa Liidu eramajapidamistes

See märk näitab, et seadet ei tohi visata olmeprügi hulka. Inimeste tervise ja keskkonna säästmise nimel tuleb äravisatav toode tuua elektriliste ja elektrooniliste seadmete käitlemisega egelevasse kogumispunkti. Küsimuste korral pöörduge kohaliku prügikäitlusettevõtte poole.

Finnish notice

Kotitalousjätteiden hävittäminen Euroopan unionin alueella

Tämä symboli merkitsee, että laitetta ei saa hävittää muiden kotitalousjätteiden mukana. Sen sijaan sinun on suojattava ihmisten terveyttä ja ympäristöä toimittamalla käytöstä poistettu laite sähkö- tai elektroniikkajätteen kierrätyspisteeseen. Lisätietoja saat jätehuoltoyhtiöltä.

French notice

Mise au rebut d'équipement par les utilisateurs privés dans l'Union Européenne

Ce symbole indique que vous ne devez pas jeter votre produit avec les ordures ménagères. Il est de votre responsabilité de protéger la santé et l'environnement et de vous débarrasser de votre équipement en le remettant à une déchetterie effectuant le recyclage des équipements électriques et électroniques. Pour de plus amples informations, prenez contact avec votre service d'élimination des ordures ménagères.

German notice

Entsorgung von Altgeräten von Benutzern in privaten Haushalten in der EU

Dieses Symbol besagt, dass dieses Produkt nicht mit dem Haushaltsmüll entsorgt werden darf. Zum Schutze der Gesundheit und der Umwelt sollten Sie stattdessen Ihre Altgeräte zur Entsorgung einer dafür vorgesehenen Recyclingstelle für elektrische und elektronische Geräte übergeben. Weitere Informationen erhalten Sie von Ihrem Entsorgungsunternehmen für Hausmüll.

Greek notice

Απόρριψη άχρηστου εξοπλισμού από ιδιώτες χρήστες στην Ευρωπαϊκή Ένωση

Αυτό το σύμβολο σημαίνει ότι δεν πρέπει να απορρίψετε το προϊόν με τα λοιπά οικιακά απορρίμματα. Αντίθετα, πρέπει να προστατέψετε την ανθρώπινη υγεία και το περιβάλλον παραδίδοντας τον άχρηστο εξοπλισμό σας σε εξουσιοδοτημένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Για περισσότερες πληροφορίες, επικοινωνήστε με την υπηρεσία απόρριψης απορριμμάτων της περιοχής σας.

Hungarian notice

A hulladék anyagok megsemmisítése az Európai Unió háztartásaiban

Ez a szimbólum azt jelzi, hogy a készüléket nem szabad a háztartási hulladékkal együtt kidobni. Ehelyett a leselejtezett berendezéseknek az elektromos vagy elektronikus hulladék átvételére kijelölt helyen történő beszolgáltatásával megóvja az emberi egészséget és a környezetet.További információt a helyi köztisztasági vállalattól kaphat.

Italian notice

Smaltimento di apparecchiature usate da parte di utenti privati nell'Unione Europea

Questo simbolo avvisa di non smaltire il prodotto con i normali rifi uti domestici. Rispettare la salute umana e l'ambiente conferendo l'apparecchiatura dismessa a un centro di raccolta designato per il riciclo di apparecchiature elettroniche ed elettriche. Per ulteriori informazioni, rivolgersi al servizio per lo smaltimento dei rifi uti domestici.

Latvian notice

Europos Sąjungos namų ūkio vartotojų įrangos atliekų šalinimas

Šis simbolis nurodo, kad gaminio negalima išmesti kartu su kitomis buitinėmis atliekomis. Kad apsaugotumėte žmonių sveikatą ir aplinką, pasenusią nenaudojamą įrangą turite nuvežti į elektrinių ir elektroninių atliekų surinkimo punktą. Daugiau informacijos teiraukitės buitinių atliekų surinkimo tarnybos.

Lithuanian notice

Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājsaimniecībās
Šis simbols norāda, ka ierīci nedrīkst utilizēt kopā ar citiem mājsaimniecības atkritumiem. Jums jārūpējas par cilvēku veselības un vides aizsardzību, nododot lietoto aprīkojumu otrreizējai pārstrādei īpašā lietotu elektrisko un elektronisko ierīču savākšanas punktā. Lai iegūtu plašāku informāciju, lūdzu, sazinieties ar savu mājsaimniecības atkritumu likvidēšanas dienestu.

Polish notice

Utylizacja zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych
w krajach Unii Europejskiej
Ten symbol oznacza, że nie wolno wyrzucać produktu wraz z innymi domowymi odpadkami. Obowiązkiem użytkownika jest ochrona zdrowa ludzkiego i środowiska przez przekazanie zużytego sprzętu do wyznaczonego punktu zajmującego się recyklingiem odpadów powstałych ze sprzętu elektrycznego i elektronicznego. Więcej informacji można uzyskać od lokalnej firmy zajmującej wywozem nieczystości.

Portuguese notice

Descarte de equipamentos usados por utilizadores domésticos na União Europeia
Este símbolo indica que não deve descartar o seu produto juntamente com os outros lixos domiciliares. Ao invés disso, deve proteger a saúde humana e o meio ambiente levando o seu equipamento para descarte em um ponto de recolha destinado à reciclagem de resíduos de equipamentos eléctricos e electrónicos. Para obter mais informações, contacte o seu serviço de tratamento de resíduos domésticos.

Romanian notice

Casarea echipamentului uzat de către utilizatorii casnici din Uniunea Europeană
Acest simbol înseamnă să nu se arunce produsul cu alte deşeuri menajere. În schimb, trebuie să protejaţi sănătatea umană şi mediul predând echipamentul uzat la un punct de colectare desemnat pentru reciclarea echipamentelor electrice şi electronice uzate. Pentru informaţii suplimentare, vă rugăm să contactaţi serviciul de eliminare a deşeurilor menajere local.

Slovak notice

Likvidácia vyradených zariadení používateľmi v domácnostiach v Európskej únii
Tento symbol znamená, že tento produkt sa nemá likvidovať s ostatným domovým odpadom. Namiesto toho by ste mali chrániť ľudské zdravie a životné prostredie odovzdaním odpadového zariadenia na zbernom mieste, ktoré je určené na recykláciu odpadových elektrických a elektronických zariadení. Ďalšie informácie získate od spoločnosti zaoberajúcej sa likvidáciou domového odpadu.

Spanish notice

Eliminación de los equipos que ya no se utilizan en entornos domésticos de la Unión Europea
Este símbolo indica que este producto no debe eliminarse con los residuos domésticos. En lugar de ello, debe evitar causar daños a la salud de las personas y al medio ambiente llevando los equipos que no utilice a un punto de recogida designado para el reciclaje de equipos eléctricos y electrónicos que ya no se utilizan. Para obtener más información, póngase en contacto con el servicio de recogida de residuos domésticos.

Swedish notice

Hantering av elektroniskt avfall för hemanvändare inom EU
Den här symbolen innebär att du inte ska kasta din produkt i hushållsavfallet. Värna i stället om natur och miljö genom att lämna in uttjänt utrustning på anvisad insamlingsplats. Allt elektriskt och elektroniskt avfall går sedan vidare till återvinning. Kontakta ditt återvinningsföretag för mer information.

Turkish notice

Türkiye Cumhuriyeti: EEE Yönetmeli ine Uygundur

Battery replacement notices

Dutch battery notice

Verklaring betreffende de batterij

WAARSCHUWING: dit apparaat bevat mogelijk een batterij.

- Probeer de batterijen na het verwijderen niet op te laden.
- Stel de batterijen niet bloot aan water of temperaturen boven 60° C.
- De batterijen mogen niet worden beschadigd, gedemonteerd, geplet of doorboord.
- Zorg dat u geen kortsluiting veroorzaakt tussen de externe contactpunten en laat de batterijen niet in aanraking komen met water of vuur.
- Gebruik ter vervanging alleen door HP goedgekeurde batterijen.

Batterijen, accu's en accumulators mogen niet worden gedeponeerd bij het normale huishoudelijke afval. Als u de batterijen/accu's wilt inleveren voor hergebruik of op de juiste manier wilt vernietigen, kunt u gebruik maken van het openbare inzamelingssysteem voor klein chemisch afval of ze terugsturen naar HP of een geautoriseerde HP Business of Service Partner.

Neem contact op met een geautoriseerde leverancier of een Business of Service Partner voor meer informatie over het vervangen of op de juiste manier vernietigen van accu's.

French battery notice

Avis relatif aux piles

AVERTISSEMENT : cet appareil peut contenir des piles.

- N'essayez pas de recharger les piles après les avoir retirées.
- Évitez de les mettre en contact avec de l'eau ou de les soumettre à des températures supérieures à 60°C.
- N'essayez pas de démonter, d'écraser ou de percer les piles.
- N'essayez pas de court-circuiter les bornes de la pile ou de jeter cette dernière dans le feu ou l'eau.
- Remplacez les piles exclusivement par des pièces de rechange HP prévues pour ce produit.

Les piles, modules de batteries et accumulateurs ne doivent pas être jetés avec les déchets ménagers. Pour permettre leur recyclage ou leur élimination, veuillez utiliser les systèmes de collecte publique ou renvoyez-les à HP, à votre Partenaire Agréé HP ou aux agents agréés.

Contactez un Revendeur Agréé ou Mainteneur Agréé pour savoir comment remplacer et jeter vos piles.

German battery notice

Hinweise zu Batterien und Akkus

VORSICHT: Dieses Produkt enthält unter Umständen eine Batterie oder einen Akku.

- Versuchen Sie nicht, Batterien und Akkus außerhalb des Gerätes wieder aufzuladen.
- Schützen Sie Batterien und Akkus vor Feuchtigkeit und Temperaturen über 60°.
- Verwenden Sie Batterien und Akkus nicht missbräuchlich, nehmen Sie sie nicht auseinander und vermeiden Sie mechanische Beschädigungen jeglicher Art.
- Vermeiden Sie Kurzschlüsse, und setzen Sie Batterien und Akkus weder Wasser noch Feuer aus.
- Ersetzen Sie Batterien und Akkus nur durch die von HP vorgesehenen Ersatzteile.

Batterien und Akkus dürfen nicht über den normalen Hausmüll entsorgt werden. Um sie der Wiederverwertung oder dem Sondermüll zuzuführen, nutzen Sie die öffentlichen Sammelstellen, oder setzen Sie sich bezüglich der Entsorgung mit einem HP Partner in Verbindung.

Weitere Informationen zum Austausch von Batterien und Akkus oder zur sachgemäßen Entsorgung erhalten Sie bei Ihrem HP Partner oder Servicepartner.

Italian battery notice

Istruzioni per la batteria

AVVERTENZA: Questo dispositivo può contenere una batteria.

- Non tentare di ricaricare le batterie se rimosse.
- Evitare che le batterie entrino in contatto con l'acqua o siano esposte a temperature superiori a 60° C.
- Non smontare, schiacciare, forare o utilizzare in modo improprio la batteria.
- Non accorciare i contatti esterni o gettare in acqua o sul fuoco la batteria.
- Sostituire la batteria solo con i ricambi HP previsti a questo scopo.

Le batterie e gli accumulatori non devono essere smaltiti insieme ai rifiuti domestici. Per procedere al riciclaggio o al corretto smaltimento, utilizzare il sistema di raccolta pubblico dei rifiuti o restituirli a HP, ai Partner Ufficiali HP o ai relativi rappresentanti.

Per ulteriori informazioni sulla sostituzione e sullo smaltimento delle batterie, contattare un Partner Ufficiale o un Centro di assistenza autorizzato.

Japanese battery notice

バッテリに関する注意

警告:本製品はバッテリを内蔵している場合があります。

- バッテリを取り外している場合は、充電しないでください。
- バッテリを水にさらしたり、60°C (140°F) 以上の温度にさらさないでください。
- バッテリを誤用、分解、破壊したり、穴をあけたりしないでください。
- 外部極を短絡させたり、火や水に投棄しないでください。
- バッテリを交換する際は、HP指定の製品と交換してください。

バッテリ、バッテリ パック、蓄電池は一般の家庭廃棄物と一緒に廃棄しないでください。リサイクルまたは適切に廃棄するため、公共の収集システム、HP、HPパートナー、またはHPパートナーの代理店にお送りください。

バッテリ交換および適切な廃棄方法についての情報は、HPのサポート窓口にお問い合わせください。

Spanish battery notice

Declaración sobre las baterías

ADVERTENCIA: Este dispositivo podría contener una batería.

- No intente recargar las baterías si las extrae.
- Evite el contacto de las baterías con agua y no las exponga a temperaturas superiores a los 60 °C (140 °F).
- No utilice incorrectamente, ni desmonte, aplaste o pinche las baterías.
- No cortocircuite los contactos externos ni la arroje al fuego o al agua.
- Sustituya las baterías sólo por el repuesto designado por HP.

Las baterías, los paquetes de baterías y los acumuladores no se deben eliminar junto con los desperdicios generales de la casa. Con el fin de tirarlos al contenedor de reciclaje adecuado, utilice los sistemas públicos de recogida o devuélvalas a HP, un distribuidor autorizado de HP o sus agentes.

Para obtener más información sobre la sustitución de la batería o su eliminación correcta, consulte con su distribuidor o servicio técnico autorizado.

Index

A

access rights, managing, 72
ACL, defining, 55
Array Configuration Utility, 29
array controller, purpose, 21
arrays, defined, 21
audience, 11

B

backup, with shadow copies, 47
basic disks, 23
basic disks, 23, 24
battery replacement notices, 95

C

cache file, shadow copies, 38
Canadian notice, 84
CIFS, share support, 56
cluster
  adding new storage, 73
  analysis, 69
  concepts, 62
  concepts, diagram, 62
  diagram, 60
  dual data paths, 66
  geographically dispersed, 70
  group, 70
  groups, node-based, 71
  installation, 67
  installation checklist, 66
  load balancing, 71
  managing access rights, 72
  managing file share permissions, 72
  network requirements, 66
  nodes
    powering down, 77
    powering up, 78
    restarting, 77
  overview, 59, 59
  preparing for installation, 66
  printer spooler, 75
  protocols, non cluster aware, 73
  resources, 70
  resources, defined, 60
  setting up user account, 68
clustered server elements, 25
Command View EVA
  expanding storage, 33
configuring
  private network adapter, 68
  shared disks, 69
connectivity, verifying, 68
conventions
  document, 11
  text symbols, 12
customer self repair, 13

D

data blocks, 21
data striping, 21
Declaration of Conformity, 84
disk access, verifying, 69

Disk Management
  extending volumes, 33
document
  conventions, 11
  related documentation, 11
documentation
  HP website, 11
  providing feedback, 14
domain membership, verifying, 68
dual data paths, 66
dynamic disks
  clustering, 24
dynamic disks
  spanning multiple LUNs, 23

E

European Union notice, 85
expanding storage
  Array Configuration Utility, 33
  Command View EVA, 33
extending volumes
  Disk Management, 33

F

failover
  automatic, 76
  defined, 61
  resources, 61
fault tolerance, 22
Federal Communications Commission notice, 83
file share resources, 73
file level permissions, 48
file recovery, 45
file screening management, 57
File Server Resource Manager, 27, 56
file services management, 28
file share permissions, managing, 72
file share resource planning, 71
file share resources, 63
file system elements, 24
file-sharing protocols, 24
files, ownership, 53
folder management, 48
folder recovery, 45
folders
  auditing access, 51
  managing, 48

G

group, cluster, 63
groups, adding to permissions list, 49

H

help
  obtaining, 13
HP
  Array Configuration Utility, 29
  Storage Manager, 29
  technical support, 13

I

installation, cluster, preparing for, 66
IP address resource, 63

J

Japanese notices, 85

K

Korean notices, 86

L

LAN icons, renaming, 68
laser compliance notices, 87
load balancing, 71
logical storage elements, 24
logical storage elements, 22
LUNs
  described, 23
  presenting to cluster node, 70

M

mount points
  creating, 23
  not supported with NFS, 23
mounted drives and shadow copies, 37

N

network name resource, 63
network planning, 64
network requirements, cluster, 66
NFS share resource, 74
node, server, 60

O

online spares, 22

P

partitions
  extended, 23
  primary, 23
permissions
  file level, 48
  list
    adding users and groups, 49
    removing users and groups, 49
  modifying, 50
  resetting, 50
physical disk resources, 63, 73
physical storage elements, 20
planning
  network, 64
  protocol, 65
  storage, 64
printer spooler, creating in a cluster, 75
private network adapter, configuring, 68
protocols
  non cluster aware, 73
  planning, 65
public network adapter, configuring, 68

Q

Quorum disk
  defined, 61
  recommendations, 68
quota management, 56

R

rack stability
  warning, 12
RAID
  data striping, 21
  LUNs in volumes, 23
  summary of methods, 22
regulatory compliance
  Canadian notice, 84
  European Union notice, 85
  identification numbers, 83
  Japanese notices, 85
  Korean notices, 86
  Taiwanese notices, 86
regulatory compliance
  laser, 87
related documentation, 11
resources, cluster, 60

S

SAN environment, 28
security
  auditing, 51
  file level permissions, 48
  ownership of files, 53
Services for UNIX, 25
Services for UNIX, 23
shadow copies
  in a cluster, 74
  uses, 34
shadow copies, 24
  backups, 47
  cache file, 38
  defragmentation, 37
  described, 34
  disabling, 41
  file or folder recovery, 45
  managing, 37
  mounted drives, 37
  on NFS shares, 44
  on SMB shares, 43
  planning, 34
  redirecting, 41
  scheduling, 40
  viewing list, 40
Shadow Copies for Shared Folders, 42
share management, 54
shared disks, configuring, 69
shares
  administrative, 56
  managing, 54
  standard, 56
Single Instance Storage, 27
storage management
  elements, 19
  overview, 19
  process, 20
Storage Manager for SANs, 27
storage reports, 57
storage, adding to a cluster, 73
Subscriber's Choice, HP, 13
symbols in text, 12
system updates, 80

T

Taiwanese notices, 86
technical support
  HP, 13
technical support
  service locator website, 13
text symbols, 12
troubleshooting, 79

U

user account, setting up, 68
users
  adding to permission list, 49

V

verifying
  connectivity, 68
  disk access, 69
  domain membership, 68
  name resolution, 68
virtual server, defined, 61
Volume Shadow Copy Service, 34
volumes
  planning, 23
vssadmin tool, 37

W

warning
  rack stability, 12
WEBES (Web Based Enterprise Services, 79
websites
  customer self repair, 13
  HP , 13
  HP Subscriber's Choice for Business, 13
  product manuals, 11